Voting as Choosing the Most Representative Voter

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[joint work with Umberto Grandi]

Thoughts on the State of the Union

Is there a place for logic in research on social choice theory?

Yes, but:

- While people doing SCT do use formal methods, if they use logic at all, it's typically only fairly simplistic ideas from logic.
- While many people doing applied logic have an interest in modelling social interaction, they tend to focus on decisions and strategies of individual agents, but do not take the real social choice perspective, where the *mechanism* is central.

One attempt to try and resolve this mismatch:

ESSLLI Workshop on Logical Models of Group Decision Making
(Deadline: 25 April 2013)

Storyline

We want to identify good methods for aggregating information coming from different individuals: collective decision making.

- Known problem: the simple methods people usually use ("majority") can lead to paradoxical outcome.
- Known problem: more sophisticated methods ("distance-based") are computationally intractable.
- New idea: use an aggregation rule that identifies the "most representative" voter and just copies that voter's ballot.

We'll look into this for *binary aggregation with integrity constraints* and analyse two specific implementations of the general idea.

Take-home message: simple, but works surprisingly well.

Preference Aggregation

Expert 1: $\triangle \succ \bigcirc \succ \square$

Expert 2: $\bigcirc \succ \square \succ \triangle$

Expert 3: $\square \succ \triangle \succ \bigcirc$

Expert 4: $\square \succ \triangle \succ \bigcirc$

Expert 5: $\bigcirc \succ \square \succ \triangle$

?

Judgment Aggregation

	p	$p \to q$	q
Judge 1:	True	True	True
Judge 2:	True	False	False
Judge 3:	False	True	False

?

Multiple Referenda

	fund museum?	*fund school?*	*fund metro?*
Voter 1:	Yes	Yes	No
Voter 2:	Yes	No	Yes
Voter 3:	No	Yes	Yes

?

[Constraint: we have money for *at most two projects*]

Insights so far

If you use the *majority rule* to aggregate individual pieces of information, you might encounter a *paradox*.

Here, *paradox* means:

> each *individual* piece of input information does respect some *integrity constraint*, but the *collective* information does not.

Examples for such integrity constraints:

- being a strict linear order (rather than just any binary relation)
- being a consistent set of formulas (not just any set)
- respecting the budget constraint (as opposed to violating it)

So what's a better method of aggregation?

Distance-based Aggregation

How to avoid paradoxes?

$\rightarrow$ Only consider outcomes that respect the integrity constraint.
$\rightarrow$ Which one to pick?—the one "closest" to the individual inputs.

These considerations suggest the following rule:

- The (Hamming) *distance* between an individual input and the outcome is the number of "point decisions" on which they differ.
- Elect the (consistent/rational) outcome that *minimises* the sum of distances to the individual inputs! (+ break ties if needed)

For preference aggregation (with "point decisions" being pairwise rankings), this is the famous *Kemeny rule*. No rule is perfect, but many consider this one to be pretty much the best there is.

But: this is Θ_2^p-*complete* ("complete for parallel access to NP"). ☹

Taming the Complexity

Where does this complexity come from?
$\rightarrow$ We need to search through all candidate outcomes.

- there might be exponentially many of those
- for each of them, checking consistency might be nontrivial

An idea:

- restrict set of choices to a small set of candidate outcomes
- make sure you can be sure all candidate outcomes are consistent

The easiest way of doing this:

candidate outcomes = choices made by individuals ("support")

Example

Find the outcome that minimises the sum of distances for this profile:

Issue:	**1**	**2**	**3**
20 voters:	0	1	1
10 voters:	1	0	1
11 voters:	1	1	0

Solution: $(1, 1, 1)$. The distance is 41 (41 voters $\times$ 1 disagreement).

Note: same as majority outcome (as there's no integrity constraint).

Now suppose there's an IC that says that $(1, 1, 1)$ is not ok.

Example (continued)

Find the outcome that minimises the sum of distances for this profile:

Issue:	**1**	**2**	**3**
20 voters:	0	1	1
10 voters:	1	0	1
11 voters:	1	1	0

"Average voter" says: $(0, 1, 1)$.
The distance is 42 (20 with no disagreements + 21 with 2 each).

So: not much worse (42 *vs.* 41), but easier to find (choose from 3 rather than $2^3 = 8$ outcomes; all 3 known to be consistent *a priori*)

Binary Aggregation with Integrity Constraints

Basic terminology and notation:

- Set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$; set of *issues* $\mathcal{I} = \{1, \ldots, m\}$.
- Corresponding set of *propositional symbols* $PS = \{p_1, \ldots, p_m\}$ and *propositional language* $\mathcal{L}_{PS}$ interpreted on $\mathcal{D} = \{0, 1\}^m$.
- An *aggregation rule* is a function $F : \mathcal{D}^n \to \mathcal{D}$. That is, each individual $i \in \mathcal{N}$ votes by submitting a *ballot* $B_i \in \mathcal{D}$.
- An *integrity constraint* is a formula $\mathrm{IC} \in \mathcal{L}_{PS}$ encoding a "rationality assumption". Ballot $B \in \mathcal{D}$ is *rational* iff $B \models \mathrm{IC}$. [Paradox: $\langle F, \mathrm{IC}, \boldsymbol{B} \rangle$ with $B_i \models \mathrm{IC}$ for all $i \in \mathcal{N}$ but $F(\boldsymbol{B}) \not\models \mathrm{IC}$]

And some technicalities:

- *Hamming distance* between ballots: $H(B, B') = |\{j \in \mathcal{I} \mid b_j \neq b'_j\}|$ and between a ballot and a profile: $\mathcal{H}(B, \boldsymbol{B}) = \sum_{i \in \mathcal{N}} H(B, B_i)$.
- *Support* of profile $\boldsymbol{B}$: $\textsc{Supp}(\boldsymbol{B}) = \{B_1\} \cup \cdots \cup \{B_n\}$.

Rules Based on Representative Voters

Idea: Choose an outcome by first choosing a voter (based on the input profile) and then copying that voter's ballot.

Fix $g : \mathcal{D}^n \to \mathcal{N}$. Then let $F : \boldsymbol{B} \mapsto B_{g(\boldsymbol{B})}$.

Good properties (of all these rules):

- *No paradoxes* ever, whatever the IC (not true for any other rule).
- *Unanimity* guaranteed. [obvious]
- *Neutrality* guaranteed. [maybe less obvious]
- *Low complexity* for natural choices of g.

But:

- Includes some really bad rules, such as Arrovian *dictatorships:*

  $g \equiv i$, i.e., $F : (B_1, \ldots, B_n) \mapsto B_i$ with i being the dictator

Two Representative-Voter Rules

The *average-voter rule* selects those individual ballots that minimise the Hamming distance to the profile:

$$\mathrm{AVR}(\boldsymbol{B}) \quad = \quad \underset{B \in \textsc{Supp}(\boldsymbol{B})}{\operatorname{argmin}} \; \mathcal{H}(B, \boldsymbol{B})$$

Remark: if you replace the set $\textsc{Supp}(\boldsymbol{B})$ by $\mathrm{Mod}(\mathrm{IC})$, the set of *all* consistent outcomes, you obtain the full distance-based rule.

The *majority-voter rule* selects those individual ballots that minimise the Hamming distance to one of the majority outcomes:

$$\mathrm{MVR}(\boldsymbol{B}) \quad = \quad \underset{B \in \textsc{Supp}(\boldsymbol{B})}{\operatorname{argmin}} \; \min\{H(B, B') \mid B' \in \mathrm{Maj}(\boldsymbol{B})\}$$

Connections:

- AVR related to Kemeny rule in voting/preference aggregation.
- MVR related to Slater rule in voting/preference aggregation.

Example

The AVR and the MVR really can give different outcomes:

Issue:	**1**	**2**	**3**	**4**	**5**
1 voter:	0	1	1	1	1
2 voters:	1	0	0	0	0
10 voters:	0	1	1	0	0
10 voters:	0	0	0	1	1
Maj:	0	0	0	0	0
MVR:	1	0	0	0	0
AVR:	0	0	0	1	1

Which rule is better?

We will compare the AVR and the MVR according to

- algorithmic efficiency [MVR wins]
- satisfaction of a choice-theoretic axiom [AVR wins]
- relative distance to the input profile [AVR wins]

Algorithmic Efficiency

Recall: m is the number of issues; n is the number of voters.

Winner determination for the MVR is in $O(mn)$:

- compute the majority vector in $O(mn)$
- compare each ballot to the majority vector in $O(mn)$

Winner determination for the AVR is in $O(mn \log n)$:

- compute the vector of sums in $O(mn)$
- compute the difference between each ballot (multiplied by n) to the vector of sums in $O(mn \log n)$
  [$O(\log n)$ because we are working with integers up to n]

So: both rules are efficient, but the MVR more so.

Axiom: Reinforcement

We are looking for an axiom that separates the two rules . . .

F satisfies *reinforcement* if for any two profiles $\boldsymbol{B}$ and $\boldsymbol{B}'$ with

- $\textsc{Supp}(\boldsymbol{B}) = \textsc{Supp}(\boldsymbol{B}')$ and
- $F(\boldsymbol{B}) \cap F(\boldsymbol{B}') \neq \emptyset$

it is the case that $F(\boldsymbol{B} \oplus \boldsymbol{B}') = F(\boldsymbol{B}) \cap F(\boldsymbol{B}')$.

This is a natural requirement: if two groups independently agree that a certain outcome is best, we would expect them to uphold this choice when choosing together.

Theorem: The AVR satisfies reinforcement, but the MVR does not.

Relative Distance to the Input Profile

Both rules select from $\textsc{Supp}(\boldsymbol{B})$ and the AVR by definition picks the candidate outcome closest to the profile. Thus:

Fact: The Hamming distance between the (worst) AVR-winner and the profile never exceeds the Hamming distance between the (best) MVR-winner and the profile.

More importantly: both rules are very good approximations of the full distance-based rule . . .

Approximation Results

F is said to be an *α-approximation* of F' if for every profile $\boldsymbol{B}$:

$$\max \mathcal{H}(F(\boldsymbol{B}), \boldsymbol{B}) \leqslant \alpha \cdot \min \mathcal{H}(F'(\boldsymbol{B}), \boldsymbol{B})$$

If F' is a "nice" but computationally intractable rule and if α is a constant, then this would be considered great news for F.

Theorem: Both the AVR and the MVR are (strict) *2-approximations* of the full *distance-based rule* (for any IC).

One of the insights here is that approximations get better as we increase the logical strength of the IC.

Last Slide

This work is part of an effort to better understand the very powerful framework of *binary aggregation with integrity constraints*. The focus today has been on identifying *good and simple rules to use in practice*.

- Simple (maybe simplistic) idea: pick a representative voter + copy
- Surprisingly: this works very well; we get good properties:
  - guarantee to never encounter a paradox
  - low complexity
  - good social choice-theoretic axioms (though not independence)
  - good approximation ratios wrt. distance-based rule
- Future work:
  - investigate Tideman-type rule: fix issues in order of support
  - investigate more specific aggregation frameworks and maybe get sharper approximation ratios