Voting as Selection of the Most Representative Voter

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

[joint work with Umberto Grandi (Toulouse)]

Computational Social Choice

Social choice theory deals with the aggregation of information coming from different individual agents, for collective decion making:

- voting and preference aggregation
- fair allocation of resources
- matching and coalition formation
- judgment aggregation

Traditionally studied in economics (and political science, philosophy, and mathematics), but now also in computer science and AI:

- applications: multiagent sys, recommender sys, crowdsourcing, ...
- new models: preferences, fairness, ...
- CS: algorithms and complexity, approximation, communication
- AI: knowledge representation and reasoning, machine learning

F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. Cambridge University Press, 2015. In press.

Outline

- Examples
- Binary Aggregation with Integrity Constraints
- Representative-Voter Rules
- Approximation Results

U. Grandi and U. Endriss. Lifting Integrity Constraints in Binary Aggregation. *Artificial Intelligence*, 199–200:45–66, 2013.

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. Proc. AAAI-2014.

Preference/Rank 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*]

General Perspective

The last example is actually pretty general. We can rephrase many aggregation problems as problems of *binary aggregation:*

Do you rank option $\triangle$ *above option* $\bigcirc$*?*	Yes/No
Do you believe formula "$p \rightarrow q$*" is true?*	Yes/No
Do you want the new school to get funded?	Yes/No

Each problem domain comes with its own *rationality constraints:*

Rankings should be transitive and not have any cycles.

The accepted set of formulas should be logically consistent.

We should fund at most two projects.

The *paradoxes* we have seen show that the *majority rule* does not *lift* our rationality constraints from the *individual* to the *collective* level.

Binary Aggregation with Integrity Constraints

The model:

- Set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$. Set of *issues* $\mathcal{I} = \{1, \ldots, m\}$.
- *Integrity constraint* IC: propositional formula over $\{p_1, \ldots, p_m\}$.
- *Ballot* $B \in \{0,1\}^m$ *rational* if $B \models \mathrm{IC}$. *Profile* $\boldsymbol{B} = (B_1, \ldots, B_n)$.
- *Aggregator* $F : (\{0,1\}^m)^n \to \{0,1\}^m$. Would like $F(\boldsymbol{B}) \models \mathrm{IC}$.

Example:

- $\mathcal{N} = \{1,2,3\}$. $\mathcal{I} = \{\mathtt{mus}, \mathtt{sch}, \mathtt{met}\}$. $\mathrm{IC} = \neg(\mathtt{mus} \wedge \mathtt{sch} \wedge \mathtt{met})$.
- Profile: $\boldsymbol{B} = (B_1, B_2, B_3)$ with

$$\begin{aligned} B_1 &= (1,1,0) \\ B_2 &= (1,0,1) \\ B_3 &= (0,1,1) \end{aligned}$$

$B_i \models \mathrm{IC}$ for all $i \in \mathcal{N}$, but $\mathrm{Maj}(\boldsymbol{B}) = (1,1,1)$ and $(1,1,1) \not\models \mathrm{IC}$.

Distance-based Aggregation

How to avoid paradoxes?

→ Only consider outcomes that respect the integrity constraint.
→ 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 issues on which they differ.
- Elect the rational outcome that *minimises* the sum of distances to the individual inputs! (+ break ties if needed)

For rank aggregation (with issues being pairwise rankings), this is the *Kemeny rule* (widely considered a pretty good choice).

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

Taming the Complexity

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

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

An idea:

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

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 rational *a priori*)

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 $g : (\{0,1\}^m)^n \to \mathcal{N}$. 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

Additional Notation and Terminology

- *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 $\boldsymbol{B}$: $\textsc{Supp}(\boldsymbol{B}) = \{B_1, \ldots, B_n\}$.

Two Representative-Voter Rules

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

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

Remark: if you replace the set $\textsc{Supp}(\boldsymbol{B})$ by $\mathrm{Mod}(\mathrm{IC})$, the set of *all* rational 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:

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

Connections:

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

Example

The AVR and the MVR really can give different outcomes:

Issue:	**1**	**2**	**3**	**4**	**5**	**6**
1 voter:	1	0	0	0	0	0
10 voters:	0	1	1	0	0	0
10 voters:	0	0	0	1	1	1
Maj:	0	0	0	0	0	0
MVR:	1	0	0	0	0	0
AVR:	0	1	1	0	0	0

Two More Representative-Voter Rules

We can also adapt *Tideman*'s ranked-pairs rule from voting theory. The *ranked-voter rule* (RVR) works as follows:

- order the issues by majority strength
- lock in issues in order of majority strength, whilst ensuring that the outcome remains within the support

The *plurality-voter rule* (PVR) selects the ballot chosen most often:

$$\mathrm{PVR}(\boldsymbol{B}) = \underset{B \in \textsc{Supp}(\boldsymbol{B})}{\operatorname{argmax}} |\{i \in \mathcal{N} \mid B = B_i\}|$$

The rank aggregation version of this rule has recently been proposed as a good maximum likelihood estimator by Caragiannis, Procaccia, and Shah ("modal ranking rule").

Approximation

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

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

How well do our rules F approximate the *distance-based rule* F'?

- AVR: average-voter rule
- MVR: majority-voter rule
- RVR: ranked-voter rule
- PVR: plurality-voter rule
- Arrovian dictatorships $F_i : \boldsymbol{B} \mapsto B_i$

Good would be: α is a (small) *constant*

Bad would be: α depends on n or m, not bounded by any constant

Focus on $\mathrm{Maj} = \mathrm{DBR}^{\top}$: harder to approximate than any other $\mathrm{DBR}^{\mathrm{IC}}$.

Very bad: Dictatorships

What's the worst possible scenario?

- one voter says $111\cdots111$, all others $(n{-}1)$ say $000\cdots000$
- majority rule would pick $000\cdots000$: *distance* m
- your rule picks $111\cdots111$: distance $m \cdot (n{-}1)$

Thus: worst approx. ratio for any rep-voter rule is $\frac{m \cdot (n-1)}{m} \in O(n)$

Arrovian dictatorships are maximally bad (unsurprisingly):

Proposition 1 *Every Arrovian dictatorship $F_i : \boldsymbol{B} \mapsto B_i$ is a $\Theta(n)$-approximation of the majority rule.*

Proof: See above example, with dictator saying $111\cdots111$. ✓

Almost as bad (!): RVR and PVR

Recall two of our more sophisticated rules:

- RVR: fix issues by majority strength, staying within support
- PVR: return most frequent ballot

Bad news:

Theorem 2 RVR *and* PVR *are* $\Theta(n)$*-approximations of* Maj*.*

Proof idea:

$$
\begin{array}{ll}
 & \overbrace{0\,1\,1\,1\,1\cdots1\,1}^{n-2}\,\overbrace{1\cdots\cdots\cdots1}^{m-(n-2)} \\
\text{Voter 1:} & 0\,1\,1\,1\,1\cdots1\,1\,1\cdots\cdots\cdots1 \\
\text{Voter 2:} & 1\,0\,1\,1\,1\cdots1\,1\,1\cdots\cdots\cdots1 \\
\vdots & \vdots \qquad \vdots \\
\text{Voter } n-2\text{:} & 1\,1\,1\,1\,1\cdots1\,0\,1\cdots\cdots\cdots1 \\
\text{Voter } n-1\text{:} & 1\,1\,1\,1\,1\cdots1\,1\,0\cdots\cdots\cdots0 \\
\text{Voter } n\text{:} & 1\,1\,1\,1\,1\cdots1\,1\,0\cdots\cdots\cdots0
\end{array}
$$

Remark: Similar result when assuming $m < n$, namely $\Omega(m)$.

Good: MVR and AVR

Recall: the MVR selects the ballot closest to the majority outcome.

Theorem 3 *The* MVR *is a (strict)* 2*-approximations of* Maj.

Proof idea: use triangle inequality! ✓

Recall: the AVR selects the ballot closest to the input profile. Thus:

Lemma 4 *The* AVR *approximimates* Maj *at least as well as any other representative-voter rule (thus: also a strict 2-approximation).*

Our most positive result:

Theorem 5 *Suppose* m *(the number of issues) is constant.*
Then the AVR *is a* $2\frac{m-1}{m}$*-approximation of* Maj. [*not* true for MVR]

Recall that we can get better approximation ratios for $\mathrm{IC} \neq \top$.

Other Criteria for Comparison

Complexity: Both ok, but the MVR can be computed more efficiently.

- Winner determination for the MVR is in $O(mn)$.
- Winner determination for the AVR is in $O(mn\log n)$.

Axiomatics: AVR satisfies and MVR fails a form of *reinforcement*.

$$\textsc{Supp}(\boldsymbol{B}) = \textsc{Supp}(\boldsymbol{B}') \text{ and } F(\boldsymbol{B}) \cap F(\boldsymbol{B}') \neq \emptyset \quad \Rightarrow \\ F(\boldsymbol{B} \oplus \boldsymbol{B}') = F(\boldsymbol{B}) \cap F(\boldsymbol{B}')$$

Last Slide

This work is part of a larger effort to better understand the 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 *can* work very well; we *can* get good properties:
  - guarantee to never encounter a paradox
  - low complexity
  - good social choice-theoretic axioms (though not independence)
  - for some: good approximation ratios w.r.t. distance-based rule

U. Endriss and U. Grandi. Binary Aggregation by Selection of the Most Representative Voter. Proc. AAAI-2014.