Preference and Graph Aggregation

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

[Mini Course on the Theory of Aggregation (Lecture 1)
LIP6, Pierre & Marie Curie University, Paris]

Collective Decision Making

Social choice theory is the philosophical and mathematical study of methods for *collective decision making*.

Classically, this is mostly about political decision making. But in fact the basic principles are relevant to a diverse range of questions:

- How to divide a cake between several children?
- How to assign bandwidth to competing processes on a network?
- How to choose a president given people's preferences?
- How to combine the website rankings of multiple search engines?
- How to assign student doctors to hospitals?
- How to aggregate the views of different judges in a court case?
- How to extract information from noisy crowdsourced data?

Computational social choice emphasises the fact that any method of decision making is ultimately an *algorithm*.

Example

What would be a good compromise representing the preferences of this group of five agents over three alternatives?

$$\begin{array}{rccccc}
\textbf{Agent 1:} & \triangle & \succ & \bigcirc & \succ & \square \\
\textbf{Agent 2:} & \bigcirc & \succ & \square & \succ & \triangle \\
\textbf{Agent 3:} & \square & \succ & \triangle & \succ & \bigcirc \\
\textbf{Agent 4:} & \square & \succ & \triangle & \succ & \bigcirc \\
\textbf{Agent 5:} & \bigcirc & \succ & \square & \succ & \triangle \\
\hline
 & & & \textbf{?} & &
\end{array}$$

Plan for Today

This will be an introduction to classical results (from the 1950s) on *preference aggregation*, followed by a discussion of some recent generalisations to *graph aggregation:*

- Examples for voting rules (i.e., preference aggregation rules)
- Axiomatic method: systematic study of properties of rules
- Classical results: May's Theorem and Arrow's Theorem
- Graph aggregation: framework, im/possibility results, applications

These slides are available online:

`https://staff.science.uva.nl/u.endriss/teaching/paris-2016/`

Most of the material is covered in the two papers cited below.

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

U. Endriss and U. Grandi. Collective Rationality in Graph Aggregation. Proc. 21st European Conference on Artificial Intelligence (ECAI), 2014.

Three Voting Rules

In voting, n *voters* choose from a set of m *alternatives* by stating their preferences in the form of *linear orders* over the alternatives.

Here are three *voting rules* (there are many more):

- *Plurality:* elect the alternative ranked first most often (i.e., each voter assigns 1 point to an alternative of her choice, and the alternative receiving the most points wins)
- *Plurality with runoff:* run a plurality election and retain the two front-runners; then run a majority contest between them
- *Borda:* each voter gives $m-1$ points to the alternative she ranks first, $m-2$ to the alternative she ranks second, etc.; and the alternative with the most points wins

Example: Choosing a Beverage for Lunch

Consider this election with nine *voters* having to choose from three *alternatives* (namely what beverage to order for a common lunch):

2 Germans:	Beer	$\succ$	Wine	$\succ$	Milk
3 Frenchmen:	Wine	$\succ$	Beer	$\succ$	Milk
4 Dutchmen:	Milk	$\succ$	Beer	$\succ$	Wine

Which beverage *wins* the election for

- the plurality rule?
- plurality with runoff?
- the Borda rule?

Axiomatic Method

So how do you decide which is the right voting rule to use?

The classical approach is to use the *axiomatic method:*

- identify good axioms: normatively appealing high-level properties
- give mathematically rigorous definitions of these axioms
- explore the consequences of the axioms

The definitions on the following slide are only sketched, but can be made mathematically precise (see the paper cited below for how).

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*. College Publications, 2011.

May's Theorem

When there are only *two alternatives*, then all the voting rules we have seen coincide. This is usually called the *simple majority rule* (SMR).

Intuitively, it does the "right" thing. Can we make this precise? *Yes!*

Theorem 1 (May, 1952) *A voting rule for two alternatives satisfies anonymity, neutrality, and positive responsiveness iff it is the SMR.*

Meaning of these *axioms:*

- *anonymity* = voters are treated symmetrically
- *neutrality* = alternatives are treated symmetrically
- *positive responsiveness* = if x is the (sole or tied) winner and one voter switches from y to x, then x becomes the sole winner

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

Proof Sketch

We want to prove:

> *A voting rule for two alternatives satisfies anonymity, neutrality, and positive responsiveness iff it is the SMR.*

Proof: Clearly, the simple majority rule has all three properties. ✓

Other direction: assume *#voters* is *odd* (other case: similar) $\leadsto$ no ties

Let a $\mathcal{X}$ be the set of voters voting $x \succ y$ and $\mathcal{Y}$ those voting $y \succ x$. *Anonymity* $\leadsto$ only number of ballots of each type matters. Two cases:

- If $|\mathcal{X}| = |\mathcal{Y}| + 1$, then *only x wins*. Then, by *PR*, only x wins whenever $|\mathcal{X}| > |\mathcal{Y}|$ and thus, by *neutrality*, only y wins whenever $|\mathcal{Y}| > |\mathcal{X}|$ (which is exactly the simple majority rule). ✓
- There exist $\mathcal{X}$, $\mathcal{Y}$ with $|\mathcal{X}| = |\mathcal{Y}| + 1$ but *y wins*. Let one x-voter switch to y. By *PR*, now only y wins. But now $|\mathcal{Y}'| = |\mathcal{X}'| + 1$, which is symmetric to the first situation, so by *neutrality* x wins. ↯

Condorcet Paradox

Our initial example showed that for *three or more alternatives*, the *simple majority rule* sometimes produces a *cycle*. Simpler example:

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

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

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

This is known as the *Condorcet Paradox*. Is there a better rule?

Preference Aggregation

A group of n *agents* express their *preferences* by each ranking a set of m *alternatives*. An *aggregation rule* F maps any such *profile* of individual preference orders to a single compromise preference order.

Two *axioms* you may want to impose on aggregation rules F:

- *Pareto* condition: if all agents rank x above y in the input profile, then so should the output order returned by F.
- *Independence of irrelevant alternatives* (IIA): the relative ranking of x and y in the output order returned by F should only depend on the relative rankings of x and y in the input profile.

Both axioms apply to all alternatives x and y.

Arrow's Theorem

Unfortunately, our requirements are too demanding:

Theorem 2 (Arrow, 1951) *Any aggregation rule for $\geqslant 3$ alternatives that satisfies the Pareto condition and IIA must be a dictatorship.*

An aggregation rule F is *dictatorship* if F always simply copies the preference order of some fixed dictator (one of the agents).

Remarks:

- Not true for 2 alternatives. Opposite direction also holds.
- Dictatorial does *not* just mean: outcome = someone's preference.

Next: Proof (following Geanakoplos, 2005).

K.J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

J. Geanakoplos. Three Brief Proofs of Arrow's Impossibility Theorem. *Economic Theory*, 26(1):211–215, 2005.

Extremal Lemma

Assume there are $\geqslant 3$ alternatives and F satisfies Pareto and IIA.

Claim: If all agents rank y either top or bottom, then so does F.

Proof: Suppose otherwise, i.e., all agents rank alternative y either top or bottom, but F does not. Write $\succ$ for the order returned by F.

(1) Then there exist alternatives x and z such that $x \succ y$ and $y \succ z$.

(2) By *IIA*, this does not change when we *move z above x* in every individual order (as doing so we don't cross the extremal y).

(3) By *Pareto*, in the new profile we must have $z \succ x$.

(4) But we still have $x \succ y$ and $y \succ z$, so by *transitivity* we get $x \succ z$.

Contradiction. ✓

Existence of an Extremal-Pivotal Agent

Fix some alternative y. Call an agent *extremal-pivotal* if she can push y from the bottom to the top in the output for at least one profile.

Claim: There exists an extremal-pivotal agent.

Proof: Consider a profile where every agent ranks y *at the bottom*. By Pareto, so does F. Let agents *switch* y *to the top*, one by one. By the Extremal Lemma, after each step, y *is still extremal* in F. By Pareto, at the end of this process, F ranks y *at the top*.

So there must be a point where y jumps from the bottom to the top. The agent making the corresponding switch is extremal-pivotal. ✓

Let $Prof_y$ the profile just before the jump and let $Prof^y$ be the profile just after the jump. Let i be the extremal-pivotal agent we found.

Dictatorship: Part 1

Recall: i is extremal-pivotal for y in $Prof_y$ (y at bottom for $\succ_i$ and $\succ$), from where she can force $Prof^y$ (y at top for $\succ_i$ and $\succ$).

Claim: Agent *i can dictate* the relative ranking under F of any two alternatives *x and z* that are different from y.

Proof: Suppose i wants to place x above z.

Let i vote as in $Prof^y$, except that she puts x at the top: $x \succ_i y \succ_i z$. Let all others rank *y as in $Prof^y$*, but otherwise vote as they please. Consider the resulting profile *$Prof$*:

- Note that in $Prof$ all relative rankings of x and y are as in $Prof_y$. So by *IIA*, we must still have $x \succ y$.
- Note that in $Prof$ all relative rankings of y and z are as in $Prof^y$. So by *IIA*, we must still have $y \succ z$.

By *transitivity*, we get $x \succ z$. By *IIA*, this continues to hold, if others change their relative rankings of alternatives other than x and z. ✓

Dictatorship: Part 2

Let i still be the extremal-pivotal agent relative to alternative y.

Claim: Agent i can also dictate the relative ranking under F of y and any other alternative x.

Proof: We can use a similar construction as before to show that for some alternative z, there must be an *agent* j that can *dictate* the relative ranking of x *and* y (both different from z).

But in profile $Prof_y$, i can dictate the relative ranking of x and y.
As there can be *at most one dictator* in any situation, we get $i = j$. ✓

Thus, agent i in fact is a dictator for any two alternatives, i.e., F is dictatorial. This proves Arrow's Theorem.

Graph Aggregation

Preferences orders are special types of graps. Let's generalise!

Fix a finite set of *vertices* V. A (directed) *graph* $G = \langle V, E \rangle$ based on V is defined by a set of *edges* $E \subseteq V \times V$ (thus: graph = edge-set).

Everyone in a finite group of *agents* $\mathcal{N} = \{1, \ldots, n\}$ provides a graph, giving rise to a *profile* $\boldsymbol{E} = (E_1, \ldots, E_n)$.

An *aggregation rule* is a function mapping profiles to collective graphs:

$$F : (2^{V \times V})^n \to 2^{V \times V}$$

Examples for aggregation rules:

- *majority rule:* accept an edge *iff* $> \frac{n}{2}$ of the agents do
- *intersection rule:* return $E_1 \cap \cdots \cap E_n$

U. Endriss and U. Grandi. Collective Rationality in Graph Aggregation. Proc. 21st European Conference on Artificial Intelligence (ECAI), 2014.

Applications

You may need to use graph aggregation in some of these situations:

- *Elections:* aggregation of preference relations
- *Consensus clustering:* aggregating outputs (equivalence classes) generated by different clustering algorithms
- Aggregation of Dungian *abstract argumentation frameworks* (graphs of attack relations between arguments)
- *Social network analysis:* aggregating influence networks
- *Epistemology:* aggregating Kripke frames for epistemic logics
  - aggregation by intersection = distributed knowledge
  - aggregation by union = shared knowledge
  - aggregation by transitive closure of union = common knowledge

Collective Rationality

Examples for typical properties a graph may or may not possess:

Reflexivity	$\forall x.xEx$
Symmetry	$\forall xy.(xEy \rightarrow yEx)$
Transitivity	$\forall xyz.(xEy \wedge yEz \rightarrow xEz)$
Seriality	$\forall x.\exists y.xEy$
Completeness	$\forall xy.[x \neq y \rightarrow (xEy \vee yEx)]$
Connectedness	$\forall xyz.[xEy \wedge xEz \rightarrow (yEz \vee zEy)]$

Aggregation rule F is *collectively rational* (CR) for graph property P if, whenever all individual graphs E_i satisfy P, so does the outcome $F(\boldsymbol{E})$.

Example: Condorcet Paradox = majority rule *not* CR for transitivity

▶ Which aggregation rules are CR for which graph properties?

Example

Three agents each provide a graph on the same set of four vertices:

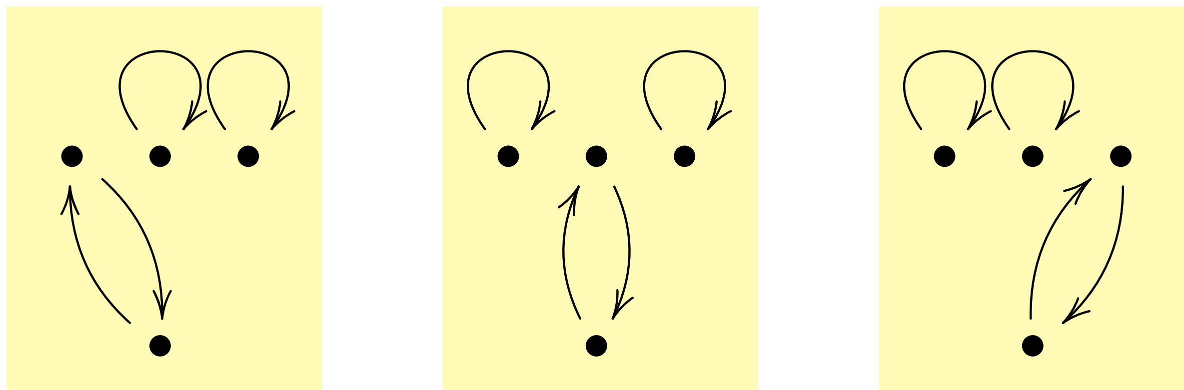

1 2 3

If we aggregate using the *majority rule*, we obtain this graph:

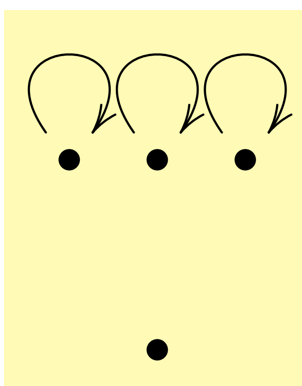

Observations:

- Majority rule not collectively rational for *seriality*.
- But *symmetry* is preserved.
- So is *reflexivity* (easy: agents violate it).

Axioms

Want to study collective rationality for *classes* of aggregation rules rather than *specific* rules (such as the majority rule).

We may want to impose certain *axioms* on $F : (2^{V \times V})^n \to 2^{V \times V}$, e.g.:

- *Anonymous:* $F(E_1, \ldots, E_n) = F(E_{\sigma(1)}, \ldots, E_{\sigma(n)})$
- *Nondictatorial:* for no $i^\star \in \mathcal{N}$ you always get $F(\boldsymbol{E}) = E_{i^\star}$
- *Unanimous:* $F(\boldsymbol{E}) \supseteq E_1 \cap \cdots \cap E_n$
- *Grounded:* $F(\boldsymbol{E}) \subseteq E_1 \cup \cdots \cup E_n$
- *Neutral:* $N_e^{\boldsymbol{E}} = N_{e'}^{\boldsymbol{E}}$ implies $e \in F(\boldsymbol{E}) \Leftrightarrow e' \in F(\boldsymbol{E})$
- *Independent:* $N_e^{\boldsymbol{E}} = N_e^{\boldsymbol{E'}}$ implies $e \in F(\boldsymbol{E}) \Leftrightarrow e \in F(\boldsymbol{E'})$

For technical reasons, we'll restrict some axioms to *nonreflexive edges*.

Notation: $N_e^{\boldsymbol{E}} = \{i \in \mathcal{N} \mid e \in E_i\}$ = *coalition* accepting edge e in $\boldsymbol{E}$

Basic Results

Proposition 3 *Every unanimous aggregation rule is CR for reflexivity.*

Proof: If every individual graph includes edge (x, x), then unanimity ensures the same for the collective outcome graph. ✓

Proposition 4 *Every grounded aggregation rule is CR for irreflexivity.*

Proof: Similar. ✓

Proposition 5 *Every neutral aggregation rule is CR for symmetry.*

Proof: If (x, y) and (y, x) have the same support, neutrality ensures that either both or neither are accepted. ✓

General Impossibility Theorem

Terminology: *Arrovian* = independent + unanimous + grounded

Theorem 6 (Endriss and Grandi, 2014) *Any Arrovian aggregation rule for $\geqslant 3$ vertices that is CR for some contagious, implicative and disjunctive graph property must be dictatorial on nonreflexive edges.*

Sketchy definition of the meta-properties of graphs used here:

- Implicative $\approx [\bigwedge S^+ \wedge \neg \bigvee S^-] \to [e_1 \wedge e_2 \to e_3]$
- Disjunctive $\approx [\bigwedge S^+ \wedge \neg \bigvee S^-] \to [e_1 \vee e_2]$
- Contagious $\approx$ for every accepted edge, there are some conditions under which also one of its "neighbouring" edges is accepted

Example:

$$\left.\begin{array}{l}\bullet\ \textit{Transitivity}\text{ is contagious and implicative}\\ \bullet\ \textit{Completeness}\text{ is disjunctive}\end{array}\right\} \Rightarrow \textit{Arrow's Theorem}$$

Application: Preference Aggregation in AI

As an immediate corollary to our theorem, we get *Arrow's Theorem* (both for strict linear orders and for weak orders).

Arrow's Thm does *not* hold for for *partial-order preferences* (popular in AI), as the *intersection rule* has all the required properties. But:

Theorem 7 (Pini et al., 2009) *Any preference aggregation rule for preorders with maximal elements for three or more alternatives that is Arrovian must be a dictatorship.*

Preorders are reflexive and transitive. Having a maximal element means that at least one alternative is as good as any other.

Proof: *Transitivity* is contagious and implicative. *Maximal element* property is a disjunctive. *Irreflexivity* of the input together with groundedness means that any NR-dictator is actually a full dictator. ✓

M.S. Pini, F. Rossi, K.B. Venable, and T. Walsh. Aggregating Partially Ordered Preferences. *Journal of Logic and Computation*, 19(3):475–502, 2009.

Application: Consensus Clustering

Clustering algorithms try to partition data points into clusters. Output is an *equivalence relation* (equivalent = in same cluster). Don't want a *trivial* clustering, where every point is its own cluster.

Consensus clustering is about finding a compromise between the solutions suggested by several algorithms: need to use aggregation.

Theorem 8 *Any aggregation rule for nontrivial equivalence relations on three or more data points that is Arrovian must be a dictatorship.*

Proof: *Transitivity* is both contagious and implicative, while the *nontriviality* condition is disjunctive (disjunction over all edges). *Reflexivity* of the input together with unanimity means that any NR-dictator is actually a full dictator. ✓

Last Slide

This has been an introduction to classical *preference aggregation* as studied in social choice theory, as well as to its generalisation in the form of *graph aggregation*. Topics covered:

- Examples for voting rules (i.e., preference aggregation rules)
- Axiomatic method: systematic study of properties of rules
- Classical results: May's Theorem and Arrow's Theorem
- Graph aggregation: framework, im/possibility results, applications

Next week we will review *judgment aggregation* (again more general) and discuss an application of *collective annotation* via crowdsourcing.

Again, the slides are available online:

`https://staff.science.uva.nl/u.endriss/teaching/paris-2016/`