Logic and Social Choice Theory

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

Social Choice Theory

SCT studies collective decision making: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

$$\triangle \succ_1 \bigcirc \succ_1 \square$$

$$\square \succ_2 \triangle \succ_2 \bigcirc$$

$$\bigcirc \succ_3 \square \succ_3 \triangle$$

$$?$$

SCT is traditionally studied in Economics and Political Science, but now also by "us": *Computational Social Choice*.

Computational Social Choice

Research can be broadly classified along two dimensions —

The kind of *social choice problem* studied, e.g.:

- electing a winner given individual preferences over candidates
- aggregating individual judgements into a collective verdict
- fairly dividing a cake given individual tastes

The kind of *computational technique* employed, e.g.:

- algorithm design to implement complex mechanisms
- complexity theory to understand limitations
- logical modelling to fully formalise intuitions
- knowledge representation techniques to compactly model problems
- adaptation for deployment in a multiagent system

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. *A Short Introduction to Computational Social Choice*. Proc. SOFSEM-2007.

Talk Outline

I will briefly introduce four research areas currently active at the ILLC in which we apply logic to social choice theory:

- Logic for the compact representation of large problem instances: social choice in combinatorial domains
- Logic for the formalisation of social choice mechanisms: from the *axiomatic method* to logics for social choice
- Logic as a basis for the verification and discovery of theorems in social choice theory: automated reasoning for social choice theory
- Logic as the object of aggregation: judgment aggregation and the computational complexity of judgment aggregation

Social Choice in Combinatorial Domains

Many social choice problems have a *combinatorial structure:*

- Elect a *committee* of k members from amongst n candidates.
- Find a good *allocation* of n indivisible goods to agents.

Seemingly small problems generate huge numbers of alternatives:

- Number of 3-member committees from 10 candidates: $\binom{10}{3} = 120$ (i.e. $120! \approx 6.7 \times 10^{198}$ possible rankings)
- Allocating 10 goods to 5 agents: $5^{10} = 9765625$ allocations and $2^{10} = 1024$ bundles for each agent to think about

We need good *languages* for representing preferences!

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Ordinal Preferences: CI-Nets

Until recently there has been no compact language for ordinal preferences that are monotonic. $\rightsquigarrow$ *Conditional Importance Networks:*

- A CI-net is a set of CI-statements of the form $S^+, S^- : S_1 \triangleright S_2$. ("if I own all the items in S^+ and none of those in S^-, then obtaining set S_1 is more important to me than obtaining set S_2")
- The *preference order* induced by a CI-net is the smallest partial order that is monotonic and satisfies all its CI-statements.

We are also using (simple fragments of) CI-nets to model *fair division:*

- Given a group of agents' individual preferences over a set of indivisible goods, can we find an allocation that is envy-free?

S. Bouveret, U. Endriss, and J. Lang. CI-Nets: A Graphical Language for Representing Ordinal, Monotonic Preferences over Sets of Goods. Proc. IJCAI-2009.

S. Bouveret, U. Endriss, and J. Lang. Fair Division under Ordinal Preferences: Computing Envy-Free Allocations of Indivisible Goods. Proc. ECAI-2010.

Cardinal Preferences: Weighted Goals

Weighted goals are a logic-based language for to compactly represent utility functions over binary combinatorial domains.

- Propositional language over PS. Want to model $u : 2^{PS} \to \mathbb{R}$.
- Formulas of $\mathcal{L}_{PS}$ represent goals. Weights represent importance.
- For each truth assignment, aggregate weights of satisfied formulas.

Results include:

- *Expressivity:* with sum aggregation, positive goals with positive weights can express all monotonic functions, and only those
- *Complexity:* social welfare maximisation is NP-hard for max aggregation, even if all weighted goals have the form $(p \wedge q, 1)$

J. Uckelman. *More than the Sum of its Parts: Compact Preference Representation over Combinatorial Domains*. PhD thesis, ILLC, University of Amsterdam, 2009.

J. Uckelman and U. Endriss. Compactly Representing Utility Functions Using Weighted Goals and the Max Aggregator. *Artif. Intell.*, 174(15):1222–1246, 2010.

Finer Analysis via Linear Logic

Weighted goals *cannot* express statements such as this:

> "getting p has value 5 to me, but getting p *twice* has value 8"

But being able to model this is important for *combinatorial auctions* and *negotiation* in multiagent systems.

Resource-sensitive logics, in particular *linear logic*, can speak about the multiplicity of items.

D. Porello and U. Endriss. Modelling Combinatorial Auctions in Linear Logic. Proc. KR-2010.

D. Porello and U. Endriss. Modelling Multilateral Negotiation in Linear Logic. Proc. ECAI-2010.

The Axiomatic Method in Social Choice Theory

Modern SCT has always made use of logic, albeit informally. The first example of the "axiomatic method" was *Arrow's Theorem* (1951):

> *Any aggregation mechanism for a finite group of individuals to rank ≥ 3 alternatives that satisfies the weak Pareto condition and independence or irrelevant alternatives must be dictatorial.*

The three axioms involved are:

- *Weak Pareto:* if all individuals prefer x to y, then so should society
- *IIA:* the social ranking of x vs. y should only depend on the individual rankings of x vs. y
- *Nondictatoriality:* the aggregator should not be a dictatorship, i.e., a function that just copies the ranking of a fixed individual

Logics for Social Choice Theory

We have shown how to model the Arrovian framework of preference aggregation (PA) in FOL. Arrow's Theorem reduces to this:

$\mathrm{T}_{\mathrm{PA}} \cup \{\mathrm{PAR}, \mathrm{IIA}, \mathrm{NDIC}\}$ *does not have a finite model.*

This is interesting for (at least) two reasons:

- It tells us something about the nature of the axioms proposed in SCT (e.g., second-order quantification is *not* needed).
- It can form the basis for the verification of results in SCT using automated theorem provers.

Related work: (new) *modal logic* (Ågotnes et al., JAAMAS 2010); *propositional logic* (Tang & Lin, AIJ 2009); *HOL* (Nipkow, JAR 2009). For the latter two the focus is on automated reasoning.

U. Grandi and U. Endriss. First-Order Logic Formalisation of Arrow's Theorem. Proc. LORI-2009.

Automated Discovery of Theorems

Another area of SCT is *ranking sets of objects:* how do you extend a preference order on objects to a preference order on sets of objects?

Example: The *Kannai-Peleg Theorem* (JET, 1984) shows that for sets $\mathcal{X}$ with $|\mathcal{X}| \geq 6$ it is impossible to extend total orders $\succeq$ on $\mathcal{X}$ to weak orders $\hat{\succeq}$ on $2^{\mathcal{X}} \setminus \{\emptyset\}$ in a manner that respects:

- *Dominance:* prefer $A \cup \{x\}$ to A whenever you prefer x to all $y \in A$, and prefer A to $A \cup \{x\}$ whenever you prefer all $y \in A$ to x
- *Independence:* weakly prefer $A \cup \{x\}$ to $B \cup \{x\}$ if you (strictly) prefer A to B and $x \notin A \cup B$

Approach to derive similar new results for this domain:

- Use model-theoretic argument to show that for axioms of certain syntactic form, impossibilities established for $|\mathcal{X}| = k$ always generalise.
- Translate small instances into propositional logic and use SAT solver.

C. Geist and U. Endriss. Automated Search for Impossibility Theorems in Social Choice Theory: Ranking Sets of Objects. *Journal of AI Research*. (astmr's)

Judgment Aggregation

Preferences are not the only structures we may wish to aggregate. JA studies the aggregation of judgments on logically inter-connected propositions. Example:

	p	$p \rightarrow q$	q
Judge 1:	Yes	Yes	Yes
Judge 2:	No	Yes	No
Judge 3:	Yes	No	No
Majority:	Yes	Yes	No

Paradox: each *individual* judgment set is *consistent*, but the *collective* judgment arrived at using the *majority rule* is not

Complexity of Judgment Aggregation

We have initiated a study of the computational complexity of JA:

- Performing aggregation is polynomial for the *premise-based procedure* and NP-complete for the *distance-based procedure*.
- *Manipulation* is NP-complete for the *premise-based procedure*.
- Deciding *safety of the agenda* (does a given agenda rule out paradoxes?) is Π_2^p-complete for the *majority rule*.

  Deciding safety of the agenda wrt. a class of procedures characterised by axioms such as *anonymity*, *neutrality*, and *independence* typically has the same complexity.

U. Endriss, U. Grandi, and D. Porello. Complexity of Judgment Aggregation: Safety of the Agenda. Proc. AAMAS-2010.

U. Endriss, U. Grandi, and D. Porello. Complexity of Winner Determination and Strategic Manipulation in Judgment Aggregation. Proc. COMSOC-2010.

Last Slide

- COMSOC is an exciting area of research bringing together ideas from mathematical economics (particularly social choice theory) and computer science (including logic).
- Examples of ongoing research at the ILLC we have reviewed here:
  - Compact Representation of Preferences
  - Logical Modelling of Social Choice Mechanisms
  - Automated Reasoning for Social Choice Theory
  - Judgment Aggregation
- COMSOC website: `http://www.illc.uva.nl/COMSOC/`