

Voting in Combinatorial Domains: What Logic and AI Have to Say (Extended Abstract)

Jérôme Lang

IRIT - Université Paul Sabatier and CNRS
31062 Toulouse Cedex (France)
lang@irit.fr

This talk is half-way between a survey and a report on ongoing research, most of which is joint work with Vincent Conitzer and Lirong Xia. The first part of the talk owes a lot to two survey papers co-authored with Yann Chevaleyre, Ulle Endriss and Nicolas Maudet [5, 6].

Social Choice Theory is the subfield of economics that aims at designing and evaluating methods for collective decision making [1]. Most works in the field focus on normative questions: the problem is generally considered to be solved when the existence (or the non-existence) of a procedure meeting some requirements has been shown; how hard it is to compute this procedure, and how it should be computed, have deserved much less attention in the Social Choice community. This is where Computer Science, and more specifically Artificial Intelligence, come into play, giving birth to a new interdisciplinary research field called *Computational Social Choice* (see [5] for an introduction to the field). *The first part of the talk will consist of a brief introduction to computational social choice.*

One of the hot topics in computational social choice is *voting on a set of alternatives that has a combinatorial structure*: in other words, the voters have to make a common decision on several possibly related issues. For instance, the inhabitants of some local community may have to make a joint decision over several related issues of local interest, such as deciding whether some new public facility such as a swimming pool or a tennis court should be built. Such elections are called *multiple referenda* [4]. Some of the voters may have preferential dependencies, for instance, they may prefer the tennis court to be built only if the swimming pool is not. Another example is when the members of an association have to elect a steering committee, composed of a president, a vice-president and a treasurer (see for instance [2]). Again, voters typically have preferential dependencies.

As soon as voters have preferential dependencies between issues, it is generally a bad idea to decompose the problem into a set of p smaller problems, each one bearing on a single issue. Doing so typically lead to “multiple election paradoxes”. Such paradoxes have been studied by a number of authors [4, 2, 8]. Consider the following example. A joint decision has to be made about whether or not to build a new swimming pool (S or \bar{S}) and a new tennis court (T or \bar{T}). Assume that the preferences of voters 1 and 2 are $S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T} \succ ST$, those of voters 3 and 4 are $\bar{S}T \succ S\bar{T} \succ \bar{S}\bar{T} \succ ST$, and those of voter 5 are $ST \succ S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T}$. The first problem is that voters 1

to 4 will feel ill at ease when asked to report their projected preference on $\{S, \bar{S}\}$ and $\{T, \bar{T}\}$. Only voter 5 knows that whatever the other voters' preferences about $\{S, \bar{S}\}$ (resp. $\{T, \bar{T}\}$), she can vote for T (resp. S) without any risk of experiencing regret. Experimental studies suggest that most voters tend to report their preferences optimistically in such situations; for instance, voters 1–2 would likely report a preference for S over \bar{S} . The second problem (the paradox itself) is that under this assumption that voters report optimistic preferences, the outcome will be ST , which is *the worst outcome for all but one voter*. A difficult question is how to design a method for voting on related issues that avoids such paradoxes.

We can list five ways of proceeding, each of which has its own pitfalls: (1) ask voters to report their entire preference relation explicitly on the set of alternatives, and then apply a fixed voting rule; (2) ask voters to report only a small part of their preference relation (for instance, their k preferred outcomes, where k is a small number) and apply a voting rule that needs this information only; (3) limit the number of possible combinations that voters may vote for, as advocated in [4]; (4) ask each voter to express her preferences as an input in some fixed compact representation language, and then apply a fixed voting rule to the profile consisting of the preference relations induced by the voters' inputs (see [9]); and finally, (5) impose a domain restriction allowing the voters to vote separately on each issue, either simultaneously or sequentially. Especially, if the voters' preferences are *separable* (which means that each voter's preference on the values of an issue is independent from the outcome on other issues), then this approach is reasonably safe, as shown in [8].

The second part of the talk will present the general problem of voting on combinatorial domains, as well as its paradoxes, and will briefly discuss the pros and cons of these five possible classes of solutions.

Then we focus on the last class of solutions, which seems to be the most promising. Its main drawback is that separability is a very demanding assumption, and is unlikely to be met in practice. Several recent papers [10, 13, 14, 12] impose a much weaker domain restriction under which sequential voting can be applied "safely": each time a voter is asked to report her preferences on an issue, these preferences do not depend on the values of the issues that have not been decided yet. Formally, this can be expressed as the following condition: there is a linear order $\mathcal{O} = \mathbf{X}_1 > \dots > \mathbf{X}_p$ on the set of issues such that the preferences of each voter on \mathbf{X}_j are preferentially independent from $\mathbf{X}_{j+1}, \dots, \mathbf{X}_p$ given $\mathbf{X}_1, \dots, \mathbf{X}_{j-1}$. If this property is satisfied, then *sequential voting rules* can be defined in the following way. Let r_1, \dots, r_p be voting rules on the domains of $\mathbf{X}_1, \dots, \mathbf{X}_p$ respectively. The *sequential composition* of r_1, \dots, r_p is defined by applying the following protocol repeatedly for $i = 1, \dots, p$: elicit voters' preferences on the domain of \mathbf{X}_i ; apply r_i to these local preferences; broadcast the outcome to the voters. Clearly, in order to compute the outcome of these sequential voting rules we do not need to know the voters' full preference relations: it suffices for each voter to express a CP-net [3], with the condition that the dependency graph of the CP-nets is *acyclic* and *common to all voters*.

Even if the domain restriction imposed by sequential voting is much weaker than separability, it still eliminates most of the profiles. This restriction is weakened in [14]:

profiles must only be compatible with *some* common order, not specified in the definition of the sequential voting rule. [12] goes much further and give a generalization sequential voting rules that does not require any domain restriction. However, the gain in generality comes with a complexity gap.

The third part of the talk will present sequential voting rules and these latter two generalizations.

Coming back to the problem that voters encounter when asked to report their projected preferences on single issues (in the example above, between $\{S, \bar{S}\}$ and $\{T, \bar{T}\}$): when expressing such preferences, voters have to lift their preferences from the level of alternatives to the level of *sets of alternatives*, or equivalently, to the level of propositional formulas (with issues as propositional variables). This suggests the existence of a strong connection between voting on combinatorial domains and *preference logics* [11, 7]. *This will be the topic of the last part of the talk.*

References

1. K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*. North-Holland, 2002.
2. J.-P. Benoit and L. Kornhauser. On the assumption of separable assembly preferences. *Social Choice and Welfare*, 16(3):429–439, 1999.
3. C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
4. S. Brams, D. Kilgour, and W. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
5. Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proc. 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM-2007)*. Springer-Verlag, 2007.
6. Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine*, 2008. To appear.
7. S. O. Hansson. Preference logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, pages 319–393. Kluwer, 2001.
8. D. Lacy and E. Niou. A problem with referenda. *Journal of Theoretical Politics*, 12(1):5–31, 2000.
9. J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.
10. J. Lang. Voting and aggregation on combinatorial domains with structured preferences. In *Proc. 20th Joint International Conference on Artificial Intelligence (IJCAI-2007)*, 2007.
11. H.H. von Wright. *The logic of preference*. Edinburgh University Press, 1963.
12. L. Xia, V. Conitzer, and J. Lang. Voting on multiattribute domains with cyclic preferential dependencies – a new family of voting rules. In *Proceedings of AAAI08*, 2008.
13. L. Xia, J. Lang, and M. Ying. Sequential voting and multiple election paradoxes. In *Proc. 11th Conference on the Theoretical Aspects of Rationality and Knowledge (TARK-2007)*, 2007.
14. L. Xia, J. Lang, and M. Ying. Strongly decomposable voting rules on multiattribute domains. In *Proceedings of AAAI07*, 2007.