Some representation and computational issues in social choice

Jérôme Lang

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

Abstract. This paper briefly considers several research issues, some of which are on-going and some others are for further research. The starting point is that many AI topics, especially those related to the ECSQARU and KR conferences, can bring a lot to the representation and the resolution of social choice problems. I surely do not claim to make an exhaustive list of problems, but I rather list *some* problems that I find important, give some relevant references and point out some potential research issues[1].

1 Introduction

For a few years, Artificial Intelligence has been taking more and more interest in collective decision making. There are two main reasons for that, leading to two different lines of research. Roughly speaking, the first one is concerned with importing concepts and procedures from social choice theory for solving questions that arise in AI application domains. This is typically the case for managing societies of autonomous agents, which calls for negotiation and voting procedures. The second line of research, which is the focus of this position paper, goes the other way round: it is concerned with importing notions and methods from AI for solving questions originally stemming from social choice.

Social choice is concerned with designing and evaluating methods of collective decision making. However, it somewhat neglects computational issues: the problem is generally considered to be solved when the existence (or the non-existence) of a procedure meeting some requirements has been shown; more precisely, knowing that the procedure *can* be computed is generally enough; now, *how hard* this computation is, and how the procedure should be implemented, have deserved less attention in the social choice community. This is where AI (and operations research, and more generally computer science) comes into play. As often when bringing together two traditions, AI probably raises more new questions pertaining to collective decision making than it solves old ones. One of

[1] Writing a short survey is a difficult task, especially because it always leads to leaving some relevant references aside. I'll maintain a long version of this paper, accessible at http://www.irit.fr/recherches/RPDMP/persos/JeromeLang/papers/ecsqaru05-long.pdf, and I'll express my gratitude to everyone who'll point to me any missing relevant reference.

the most relevant of these issues consists in considering group decision making problems when the set of alternative is finite and has a *combinatorial structure*.

This paper gives a brief overview of some research issues along this line. Section 2 starts with the crucial problem of eliciting and representing the individual's preferences on the possible alternatives. Section 3 focuses on preference aggregation, Section 4 on vote, and Section 5 on fair division. Section 6 evokes other directions deliberately ignored in this short paper.

2 Elicitation and compact representation of preference

Throughout the paper, $N = \{1, \ldots, n\}$ is the (finite) set of *agents* involved in the collective choice and X is the finite set of *alternatives* on which the decision process bears.

Any individual or collective decision making problem needs some description (at least partial) of the preferences of each of the agents involved over the possible alternatives. A *numerical* preference structure is a utility function $u : X \rightarrow \mathbb{R}$. An *ordinal* preference structure is a preorder P on X, called preference relation. $R(x, y)$ is denoted alternatively by $x \succeq y$. $\succ$ denotes strict preference ($x \succ y$ if and only if $x \succeq y$ and not $y \succeq x$) and $\sim$ denotes indifference ($x \sim y$ if and only if $x \succeq y$ and $y \succeq x$). An intermediate model between pure ordinality and pure numerical models is that of *qualitative preferences*, consisting of (qualitative) utility functions $u : X \rightarrow L$, where L is a totally ordered (yet not numerical) scale. Unlike ordinal preferences, qualitative preferences allow commensurability between uncertainty and preference scales as well as interagent comparison of preferences (see [22] for discussions on ordinality in decision making.)

The choice of a *model*, i.e. a mathematical structure, for preference, does not tell how agents' preferences are obtained from them, stored, and handled by algorithms. *Preference representation* consists in choosing a language for encoding preferences so as to spare computational resources. The choice of a language is guided by two tasks: upstream, *preference elicitation* consists in interacting with the agent so as to obtain her preferences over X, while *optimization* consists in finding nondominated alternatives from a compactly represented input.

As long as the set of alternatives has a small size, the latter problems are computationally easy. Unfortunately, in many concrete problems the set of alternatives has a *combinatorial structure*. A combinatorial domain is a Cartesian product of finite value domains for each one of a set of variables: an alternative in such a domain is a tuple of values. Clearly, the size of such domains grows exponentially with the set of variables and becomes quickly very large, which makes explicit representations and straightforward elicitation and optimization no longer reasonable. Logical or graphical *compact representation languages* allow for representing in as little space as possible a preference structure whose size would be prohibitive if it were represented explicitly. The literature on preference elicitation and representation for combinatorial domains has been growing fastly for a few years, and due to the lack of space I omit giving references here.

The criteria one can use for choosing a compact preference language include, at least, the following ones:

- *cognitive relevance*: a language should be as close as possible to the way human agents "know" their preferences and express them in natural language;
- *elicitation-friendliness*: it should be easy to design algorithms to elicit preference from an agent so as to get an output expressed in a given language;
- *expressivity*: find out the set of preference relations or utility functions that can be expressible in a given language;
- *complexity*: given an input consisting of a compactly represented preference structure in a given language, determine the computational complexity of finding a non-dominated alternative, checking whether an alternative is preferred to another one, whether an alternative is non-dominated etc.;
- *comparative succinctness*: given two languages L and L', determine whether every preference structure that can be expressed in L can also be expressed in L' without a significant (suprapolynomial) increase of size, in which case L' is said to be at least as succinct as L.

Cognitive relevance is somewhat hard to assess, due to its non-technical nature, and has been rarely studied. Complexity has been studied in [35] for logic-based languages. Expressivity and comparative succinctness have been systematically investigated in [19] for *ordinal* preference representation.

Although these languages have been designed for single agents, they can be extended to multiple agents without much difficulty; [34] and [44] are two examples of such extensions.

3 Preference aggregation

Preference aggregation, even on simple domains, raises challenging computational issues that have been recently investigated by AI researchers. Aggregating preferences consist in mapping a collection $\langle P_1, \ldots, P_n \rangle$ of preference relations (or *profiles*) into a collective preference relation P^* (which implies circumvening Arrow's impossibility theorem [2] by relaxing one of its applicability conditions.)

Now, even on simple domains, some aggregation functions raise computational difficulties. This is notably the case for Kemeny's aggregation rule, consisting in aggregating the profiles into a profile (called *Kemeny consensus*) being closest to the n profiles, with respect to a distance which, roughly speaking, is the sum, for all agents, of the numbers of pairs of alternatives on which the aggregated profile disagrees with the agent's profile. Computing a Kemeny consensus is NP-hard; [21] addresses its practical computation.

When the set of alternatives has a combinatorial structure, things get much worse. Moreover, since in that case preferences are often described in a compact representation language, aggregation should ideally operate directly on this language, without generating the individual nor the aggregated preferences explicitly. A common way of aggregating compactly represented preferences is *(logical) merging*. The common point of logic-based merging approaches is that the set of alternatives corresponds to a set of propositional worlds; the logic-based representation of agent's preferences (or beliefs) then induces a cardinal function (using ranks or distances) on worlds and aggregates these cardinal preferences. These functions are not necessarily on a numerical scale but the scale has to be common to all agents. We do not have the space to give all relevant references to logic-based merging here, but we give a few of them, which explicitly mention some social choice theoretic issues: [33, 40, 13, 39]. See also [34, 6] for preference aggregation from logically expressed preferences. .

4 Vote

Voting is one of the most popular ways of reaching common decisions. Researchers in social choice theory have studied extensively the properties of various families of voting rules, but, again, have neglected computational issues. A *voting rule* maps each collection of individual preference profiles, generally consisting of linear orders over the set of candidates, to a *nonempty subset* of the set of candidates; if the latter subset is always a singleton then the voting rule is said to be *deterministic*[2].

For a panorama of voting rules see for instance [10]. We just give here a few of them. A *positional scoring rule* is defined from a *scoring vector*, i.e. a vector $\boldsymbol{s} = (s_1, \ldots, s_m)$ of integers such that $s_1 \geq s_2 \geq \ldots \geq s_m$ and $s_1 > s_m$. Let $rank_i(x)$ be the rank of x in $\succ_i$ (1 if it is the favorite candidate for voter i, 2 if it is the second favorite etc.), then the score of x is $S(x) = \sum_{i=1}^{N} s_{rank_i(x)}$. Two well-known examples of positional scoring procedures are the *Borda* rule, defined by $s_k = m - k$ for all $k = 1, \ldots, m$, and the *plurality rule*, defined by $s_1 = 1$, and $s_k = 0$ for all $k > 1$. Moreover, a *Condorcet winner* is a candidate preferred to any other candidate by a strict majority of voters. (it is well-known that there are some profiles for which no Condorcet winner exists.) Obviously, when there exists a Condorcet winner then it is unique. A *Condorcet-consistent rule* is a voting rule electing the Condorcet winner whenever there is one.

The first question that comes to mind is whether determining the outcome of an election, for a given voting procedure, is computationally challenging (which is all the more relevant as electronic voting becomes more and more popular.)

4.1 Computing the outcome of voting rules: small domains

Most voting rules among those that are practically used are computable in linear or quadratic time in the number of candidates (and almost always linear in the number of voters); thererefore, when the number of candidates is small (which is typically the case for political elections where a single person has to be elected), computing the outcome of a voting rule does not need any sophisticated algorithm. However, a few voting rules are computationally complex. Here are three

[2] The literature of social choice theory rather makes use of the terminology "voting correspondances" and "deterministic voting rules" but for the sake of simplicity we will make use of the terminology "voting rules" in a uniform way.

of them: Dodgson's rule and Young's rule both consist in electing candidates that are closest to being a Condorcet winner: each candidate is given a score that is the smallest number of exchanges of *elementary changes* in the voters' preference orders needed to make the candidate a Condorcet winner. Whatever candidate (or candidates, in the case of a tie) has the lowest score is the winner. For Dodgson's rule, an elementary change is an exchange of adjacent candidates in a voter's preference profile, while for Young's rule it is the removal of a voter. Lastly, Kemeny's voting rule elects a candidate if and only if it is the preferred candidate in some Kemeny consensus (see Section 3). Deciding whether a given candidate is a winner for any of the latter three voting rules is a $\Delta_2^{\mathsf{P}}(O(\log n))$-complete (for Dodgson's, NP-hardness was shown in [5] and $\Delta_2^{\mathsf{P}}(O(\log n))$-completeness in [30]; $\Delta_2^{\mathsf{P}}(O(\log n))$-completeness was shown in [45] for Young's and in [31] for Kemeny's.

4.2 Computing the outcome of voting rules: combinatorial domains

Now, when the set of candidates has a combinatorial structure, even simple procedures such as plurality and Borda become hard. Consider an example where agents have to agree on a common menu to be composed of a first course dish, a main course dish, a dessert and a wine, with a choice of 6 items for each. This makes 6^4 candidates. This would not be a problem if the four items to be chosen were independent from the other ones: in this case, this vote problem over a set of 6^4 candidates would come down to four independent problems over sets of 6 candidates each, and any standard voting rule could be applied without difficulty. But things get complicated if voters express dependencies between variables, such as "I prefer white wine if one of the courses is fish and none is meat, red wine if one of the courses is meat and none is fish, and in the remaining cases I would like equally red or white wine", etc.

Obviously, the prohibitive number of candidates makes it hard, or even practically impossible, to apply voting rules in a straightforward way. The computational complexity of some voting procedures when applied to compactly represented preferences on a combinatorial set of candidates has been investigated in [35]; however this paper does not address the question of *how* the outcome can be computed in a reasonable amount of time.

When the domain is large enough, computing the outcome by first generating the whole preference relations on the combinatorial domain from their compact representation is unfeasible. A first way of coping with the problem consists in contenting oneself with an *approximation* of the outcome of the election, using incomplete and/or randomized algorithms making a possible use of heuristics. This is an open research issue.

A second way consists in decomposing the vote into *local* votes on individual variables (or small sets of variables), and gathering the results. However, as soon as variables are not preferentially independent, it is generally a bad idea: "multiple election paradoxes" [11] show that such a decomposition leads to suboptimal choices, and give real-life examples of such paradoxes, including simultaneous referenda on related issues. We give here a very simple example of such a paradox. Suppose 100 voters have to decide whether to build a swimming pool or not (S), and whether to build a tennis court or not (T). 49 voters would prefer a swimming pool and no tennis court ($S\bar{T}$), 49 voters prefer a tennis court and no swimming pool ($\bar{S}T$) and 2 voters prefer to have both (ST). Voting separately on each of the issues gives the outcome ST, although it received only 2 votes out of 100 – and it might even be the most disliked outcome by 98 of the voters (for instance because building both raises local taxes too much). Now, the latter example did not work because there is a *preferential dependence* between S and T. A simple idea then consists in exploiting *preferential independencies* between variables; this is all the more relevant as graphical languages, evoked in Section 2, are based on such structural properties. The question now is to what extent we may use these preferential independencies to *decompose* the computation of the outcome into smaller problems. However, again this does not work so easily: several well-known voting rules (such as plurality or Borda) cannot be decomposed, *even when the preferential structure is common to all voters.* Most of them fail to be decomposable *even when all variables are mutually independent for all voters.* We give below an example of this phenomenon.

Consider 7 voters, a domain with two variables $\mathbf{x}$ and $\mathbf{y}$, whose domains are respectively $\{x, \bar{x}\}$ and $\{y, \bar{y}\}$, and the following preference relations, where each agent expresses his preference relation by a CP-net [7] corresponding to the following fixed preferential structure: preference on $\mathbf{x}$ is unconditional and preference on $\mathbf{y}$ may depend on the value given to $\mathbf{x}$.

3 voters	2 voters	2 voters
$\bar{x} \succ x$	$x \succ \bar{x}$	$x \succ \bar{x}$
$x : \bar{y} \succ y$	$x : y \succ \bar{y}$	$x : \bar{y} \succ y$
$\bar{x} : y \succ \bar{y}$	$\bar{x} : \bar{y} \succ y$	$\bar{x} : y \succ \bar{y}$

For instance, the first CP-net says that the voters prefer $\bar{x}$ to x unconditionally, prefer $\bar{y}$ to y when $\mathbf{x} = x$ and y to $\bar{y}$ when $\mathbf{x} = \bar{x}$. This corresponds to the following preference relations:

3 voters	2 voters	2 voters
$\bar{x}y$	xy	$x\bar{y}$
$\bar{x}\bar{y}$	$x\bar{y}$	xy
$x\bar{y}$	$\bar{x}\bar{y}$	$\bar{x}y$
xy	$\bar{x}y$	$\bar{x}\bar{y}$

The winner for the plurality rule is $\bar{x}y$. Now, the sequential approach gives the following outcome: first, because 4 agents out of 7 unconditionally prefer x over $\bar{x}$, applying plurality (as well as any other voting rule, since all reasonable voting rules coincide with the majority rule when there are only 2 candidates) locally on $\mathbf{x}$ leads to choose $\mathbf{x} = x$. Now, given $\mathbf{x} = true$, 5 agents out of 7 prefer $\bar{y}$ to y, which leads to choose $\mathbf{y} = \bar{y}$. Thus, the sequential plurality winner is $(x, \bar{y})$ – whereas the direct plurality winner is $(\bar{x}, y)$.

Such counterexamples can be found for many other voting rules. This raises the question of finding voting rules which can be decomposed into local rules (possibly under some domain restrictions), following the preferential independence structure of the voters' profiles – which is an open issue.

4.3 Manipulation

Manipulating a voting rule consists, for a given voter or coalition of voters, in expressing an insincere preference profile so as to give more chance to a preferred candidate to be elected. Gibbard and Satterthwaite's theorem [29, 47] states that if the number of candidates is at least 3, then any nondictatorial voting procedure is manipulable for some profiles.

Consider again the example above with the 7 voters[3], and the plurality rule, whose outcome is $\bar{x}y$. The two voters whose true preference is $xy \succ x\bar{y} \succ \bar{x}\bar{y} \succ \bar{x}y$ have an interest to report an insincere preference profile with $x\bar{y}$ on top, that is, to vote for $x\bar{y}$ – in that case, the winner is $x\bar{y}$, which these two voters prefer to the winner if they express their true preferences, namely $\bar{x}y$.

Since it is theoretically not possible to make manipulation impossible, one can try to make it less efficient or more difficult. Making manipulation less efficient can consist in making as little as possible of the others' votes known to the would-be manipulating voter – which may be difficult in some contexts. Making it more *difficult to compute* is a way followed recently by [4, 3, 15, 14, 17]. The line of argumentation is that if finding a successful manipulation is extremely hard computationally, then the voters will give up trying to manipulate and express sincere preferences. Note that, for once, the higher the complexity, the better. Randomization can play a role not only in making manipulation less efficient but also more complex to compute [17].

In a logical merging context (see Section 3), [27] investigate the manipulation of merging processes in propositional logic. The notion of a manipulation is however more complex to define (and several competing notions are discussed indeed), since the outcome of the process is a full preference relation.

4.4 Incomplete knowledge and communication complexity

Given some incomplete description of the voters' preferences, is the outcome of the vote determined? If not, whose preferences are to be elicited and what is relevant so as to compute the outcome? Assume, for example, that we have 4 candidates A, B, C, D and 9 voters, 4 of which vote $C \succ D \succ A \succ B$, 2 of which vote $A \succ B \succ D \succ C$ and 2 of which vote $B \succ A \succ C \succ D$, the last vote being still unknown. If the plurality rule is chosen then the outcome is already known (the winner is C) and there is no need to elicit the last voter's profile. If the Borda rule is used then the partial scores are $A : 14$, $B : 10$, $C : 14$, $D : 10$, therefore the outcome is not determined; however, we do not need to know the totality of the last vote, but we only need to know whether the last voter prefers

[3] I thank Patrice Perny, from whom I borrowed this example.

A to C or C to A. This *vote elicitation* problem is investigated from the point of view of computational complexity in [16].

More generally, communication complexity is concerned with the amount of information to be communicated so that the outcome of the vote procedure is determined: since the outcome of a voting rule is sometimes determined even if not all votes are known, this raises the question in designing protocols for gathering the information needed so as to communicate as little info as possible [18]. For example, plurality needs only to know top ranked candidates, while plurality with run-off needs the top-ranked candidates and then, after communicating the names of two finalists to the voters, which one they prefer between these two.

5 Fair division

Resource allocation of *indivisible goods* aims at assigning, to each of a set of agents N, some items from a finite set R to each of a set of agents N, given their preferences over all possible combination of objects. For the sake of simplicity, we assume here that each resource must be given to one and only one agent[4].

In *centralized* allocation problems, the assignment is determined by a central authority to which the agents have given their preferences beforehand. As it stands, a centralized fair division problem is clearly a group decision making problem on a combinatorial domain, since the number of allocations grows exponentially with the number of resources. Since the description of a fair division problem needs the specification of the agents' preferences over the set of all possible combinations of objects, elicitation and compact representation issues are highly relevant here as well. Now, is a fair division problem a vote problem, where candidates are possible allocations? Not quite, because a usual assumption is made, stating that the primary preferences expressed by agents depends only of their share, that is, agent i is indifferent between two allocations as soon as they give her the same share. Furthermore, as seen below, some specific notions for fair division problems, such as envy-freeness, have no counterpart in terms of voting.

Two classes of criteria are considered in centralized resource allocation, namely *efficiency* and *equity* (or *fairness*). At one extremity, *combinatorial auctions* consist in finding an allocation maximizing the revenue of the seller, where this revenue is the sum, over all agents, of the price that the agent is willing to pay for the combination of objects he receives in the allocation (given that these price functions are not necessarily additive.) Combinatorial auctions are a very specific, purely utilitarianistic class of allocation problems, in which considerations such as equity and fairness are not relevant. They have received an enormous at-

[4] More generally, an object could be allocated to zero, one, or more agents of N. Even if most applications require the allocation to be *preemptive* (an object cannot be allocated to more than one agent), some problems do not require it. An example of such preemption-free problems is the exploitation of shared Earth observation satellites described in [36, 8].

tention since a few years (see [20]). Here we rather focus on allocation problems where fairness is involved – in which case we speak of *fair division*.

The weakest efficiency requirement is that allocations should not be Pareto-dominated: an allocation $\pi : N \rightarrow 2^X$ is *Pareto-efficient* if and only if there is no allocation π' such that (a) for all i, $\pi'(i) \succeq_i \pi(i)$ and (b) there exists an i such that $\pi'(i) \succ_i \pi(i)$. Pareto-efficiency is purely ordinal, unlike the *utilitarianistic* criterion, applicable only when preference are numerical, under which an allocation π is preferred to an allocation π' if and only if $\sum_{i \in N} u_i(\pi(i)) > \sum_{i \in N} u_i(\pi'(i))$.

None of the latter criteria deals with fairness or equity. The most usual way of measuring equity is *egalitarianism*, which compares allocations with respect to the *leximin* ordering which, informally, works by comparing first the utilities of the least satisfied agents, and when these utilities coincide, compares the utilities of the next least satisfied agents and so on (see for instance Chapter 1 of [41]).

The *leximin* ordering does not need preferences to be numerical but only *interpersonally comparable*, that is, expressed on common scale. A purely ordinal fairness criterion is *envy-freeness* : an allocation π is *envy-free* if and only if $\pi(i) \succeq_i \pi(j)$ holds for all i and all $j \neq i$, or in informal terms, each agent is at least as happy with his share than with any other one's share. It is well-known that there exist allocation problems for which no there exists no allocation being *both Pareto-efficient and envy-free.*

In *distributed* allocation problems, agents negotiate, communicate, exchange or trade goods, in a multilateral way. Works along this line have addressed the convergence conditions towards allocations being optimal from a social point of view, depending on the acceptability criteria used by agents when deciding whether or not to agree on a propose exchange of resources, and some constraints allowed on deals – see e.g. [46, 26, 24, 23, 12]. The notion of communication complexity is revisited in [25] and reinterpreted as the minimal (with respect to some criteria) sequence of deals between agents (where minimality is with respect to a criterion that may vary, and which takes into account the number of deals and the number of objects exchanged in deals). See [38] for a survey on these issues.

Whereas social choice theory has developed an important literature on fair division, and artificial intelligence has devoted much work on the computational aspects of combinatorial auctions, computational issues in *fair division* have only started recently to be investigated. Two works addressing envy-freeness from a computational prespective are [37], who compute approximately envy-free solutions (by first making it a graded notion, suitable to optimization), and [9] who relate the search of envy-freeness and efficient allocations to some well-known problems in knowledge representation. A more general review of complexity results for centralized allocation problems in in [8]. Complexity issues for distributed allocation problems are addressed in [24].

Clearly, many models developed in the AI community should have an impact on modelling, representing compactly and solving fair division problems. Moreover, some issues addressed for voting problems and/or combinatorial auctions, such as the computational aspects of elicitation and manipulation and the role of incomplete knowledge, are still to be investigated for fair division problems.

6 Conclusion

There are many more issues for further research than those that we have briefly evoked. Models and techniques from artificial intelligence should play an important role, for (at least) the following reasons:

- the importance of *ordinal and qualitative* models in preference aggregation, vote and fair division (no need to recall that the AI research community has contributed a lot to the study of these models.) Ordinality is perhaps even more relevant in social choice than in decision under uncertainty and multicriteria decision making, due to equity criteria and the difficulty of interpersonal comparison of preference.
- the role of *incomplete knowledge*, and the need to reason about agents' beliefs, especially in utility elicitation and communication complexity issues. Research issues include various ways of applying voting and allocation procedures under incomplete knowledge, and the study of communication protocols for these issues, which may call for multiagent models of beliefs, including mutual and common belief (see e.g. [28]). Models and algorithms for group decision under uncertainty is a promising topic as well.
- the need for *compact (logical and graphical) languages* for preference elicitation and representation and measure their *spatial efficiency.* These languages need to be extended to multiple agents (such as in [44]), and aggregation should be performed directly in the language (e.g., aggregating CP-nets into a new CP-net without generating the preference relations explicitly).
- the high complexity of the tasks involved leads to interesting algorithmic problems such as finding tractable subclasses, efficient algorithms and approximation methods,using classical AI and OR techniques.
- one more relevant issue is *sequential group decision making and planning* with multiple agents. For instance, [42] address the search for an optimal path for several agents (or criteria), with respect to an egalitarianistic aggregation policy.
- measuring and localizing inconsistency among a group of agents – especially when preferences are represented under a logical form – could be investigated by extending *inconsistency measures* (see [32]) to multiple agents.

References

1. H. Andreka, M. Ryan, and P.-Y. Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12(1):13–53, 2002.
2. K. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, 1951. revised edition 1963.
3. J.J. Bartholdi and J.B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

4. J.J. Bartholdi, C.A. Tovey, and M.A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
5. J.J. Bartholdi, C.A. Tovey, and M.A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(3):157–165, 1989.
6. S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Bipolar representation and fusion of preference in the possibilistic logic framework. In *Proceedings of KR2002*, pages 421–429, 2002.
7. C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: a tool for representing and reasoning with conditional *ceteris paribus* statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
8. S. Bouveret, H. Fargier, J. Lang, and M. Lemaître. Allocation of indivisible goods: a general model and some complexity results. In *Proceedings of AAMAS 05*, 2005. Long version available at http://www.irit.fr/recherches/RPDMP/persos/JeromeLang/papers/aig.pdf.
9. S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. In *Proceedings of IJCAI-05*, 2005.
10. S. Brams and P. Fishburn. Voting procedures. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, chapter 4. Elsevier, 2004.
11. S. Brams, D. M. Kilgour, and W. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15:211–236, 1998.
12. Y. Chevaleyre, U. Endriss, and N. Maudet. On maximal classes of utility functions for efficient one-to-one negotiation. In *Proceedings of IJCAI-2005*, 2005.
13. S. Chopra, A. Ghose, and T. Meyer. Social choice theory, belief merging, and strategy-proofness. *Int. Journal on Information Fusion*, 2005. To appear.
14. V. Conitzer, J. Lang, and T. Sandholm. How many candidates are required to make an election hard to manipulate? In *Proceedings of TARK-03*, pages 201–214, 2003.
15. V. Conitzer and T. Sandholm. Complexity of manipulating elections with few candidates. In *Proceedings of AAAI-02*, pages 314–319, 2002.
16. V. Conitzer and T. Sandholm. Vote elicitation: complexity and strategy-proofness. In *Proceedings of AAAI-02*, pages 392–397, 2002.
17. V. Conitzer and T. Sandholm. Universal voting protocols to make manipulation hard. In *Proceedings of IJCAI-03*, 2003.
18. V. Conitzer and T. Sandholm. Communication complexity of common votiong rules. In *Proceedings of the EC-05*, 2005.
19. S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proceedings of KR-2004*, pages 203–212, 2004.
20. P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005. To appear.
21. A. Davenport and J. Kalagnanam. A computational study of the Kemeny rule for preference aggregation. In *Proceedings of AAAI-04*, pages 697–702, 2004.
22. D. Dubois, H. Fargier, and P. Perny. On the limitations of ordinal approaches to decision-making. In *Proceedings of KR2002*, pages 133–146, 2002.
23. P. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.
24. P. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1-2):23–46, 2005.

25. U. Endriss and N. Maudet. On the communication complexity of multilateral trading: Extended report. *Journal of Autonomous Agents and Multiagent Systems*, 2005. To appear.
26. U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negociations over resources. In *Proceedings of AAMAS-03*, 2003.
27. P. Everaere, S. Konieczny, and P.Marquis. On merging strategy-proofness. In *Proceedings of KR-2004*, pages 357–368, 2004.
28. R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
29. A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
30. E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 system is complete for parallel access to NP. *JACM*, 44(6):806–825, 1997.
31. E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. Technical report, Jenaer Schriften zur Mathematik und Informatik, October 2003.
32. A. Hunter and S. Konieczny. *Approaches to measuring inconsistent information*, pages 189–234. SpringerLNCS 3300, 2004.
33. S. Konieczny and R. Pino Pérez. Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory. *European Journal of Operational Research*, 160(3):785–802, 2005.
34. C. Lafage and J. Lang. Logical representation of preferences for group decision making. In *Proceedings of KR2000*, pages 457–468, 2000.
35. J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1):37–71, 2004.
36. M. Lemaître, G. Verfaillie, and N. Bataille. Exploiting a common property resource under a fairness constraint: a case study. In *Proceedings of IJCAI-99*, pages 206–211, 1999.
37. R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of EC'04*, 2004.
38. Agentlink technical forum group on multiagent resource allocation. http://www.doc.ic.ac.uk/ ue/MARA/, 2005.
39. P. Maynard-Zhang and D. Lehmann. Representing and aggregating conflicting beliefs. *Journal of Artificial Intelligence Research*, 19:155–203, 2003.
40. T. Meyer, A. Ghose, and S. Chopra. Social choice, merging, and elections. In *Proceedings of ECSQARU-01*, pages 466–477, 2001.
41. H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
42. P. Perny and O. Spanjaard. On preference-based search in state space graphs. In *Proceedings of AAAI-02*, pages 751–756, 2002.
43. M. S. Pini, F. Rossi, K. Venable, and T. Walsh. Aggregating partially ordered preferences: possibility and impossibility results. In *Proceedings of TARK-05*, 2005.
44. F. Rossi, K. Venable, and T. Walsh. mCP nets: representing and reasoning with preferences of multiple agents. In *Proceedings of AAAI-04*, pages 729–734, 2004.
45. J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.
46. T. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
47. M. Satterthwaite. Strategyproofness and Arrow's conditions. *Journal of Economic Theory*, 10:187–217, 1975.