Algorithmic Decision Theory Meets Logic
— Invited Talk —

Jérôme Lang(✉)

LAMSADE, CNRS and Université Paris Dauphine, Paris, France
lang@lamsade.dauphine.fr

1 Introduction

Algorithmic decision theory can be roughly defined as the design and study of languages and methods for expressing and solving various classes of decision problems, including: decision under uncertainty, sequential decision making, multicriteria decision making, collective decision making, and strategic interactions in distributed decision making. A decision problem is specified by two main components: the *preferences* of the agent(s); and the *beliefs* the agent(s) has (have) about the initial state of the world and its evolution, and possibly about the beliefs and preferences of other agents. Computational tasks involve, among others: the construction and refinement of the problem, through learning and elicitation tasks; the search for a centralized decision (for an agent or a group of agents); the impact of selfish behaviour in decentralized, multi-agent decision contexts.

Logic in algorithmic decision theory can be useful as a *declarative representation language* for the various components of the problems, and as a *generic problem solving tool*. The combination of both allow for representing *and* solving complex decision making problems. Below I point to some research issues at the meeting point of logic and algorithmic decision theory. The list is certainly not exhaustive, and it is biased towards my own work.

2 Representing and Reasoning with Preferences

2.1 Compact Representation

Domains of solutions in algorithmic decision theory often have a *combinatorial* structure of the form $A = D_1 \times \ldots \times D_p$, where each D_i is a finite set of values associated with a variable X_i. A can for instance be the set of all alternatives to choose from in many voting contexts[1] such as multiple referenda or committee elections,

Expressing preferences on such domains by listing or ranking explicitly all alternatives or solutions is practically infeasible as soon as the number of variables is more than a few units, because it puts too much communication burden

[1] See [1] for an survey of voting in combinatorial domains.

F. Calimeri et al. (Eds.): LPNMR 2015, LNAI 9345, pp. 14–19, 2015.
DOI: 10.1007/978-3-319-23264-5_2

on the agents. The AI community has produced a considerable amount of work on *compact representation languages* for preferences, aiming at expressing and processing preferences over large combinatorial domains using as few computational resources (space and time) as possible. Many of these languages are based on logic (see [2] for a survey). The most elementary language consists in specifying dichotomous preferences via propositional formulae; extensions to nondichotomous preferences consist in associating priorities or weights with formulas, using distances between interpretations, or expressing preferences between propositional formulas using a *ceteris paribus* completion principle. Logic programming languages, especially answer set programming, are also very useful: see [3–5] for surveys and [6] for a very recent development.

2.2 Preference Logics

While compact preference representation languages primarily aim at expressing succinctly preferences over combinatorial domains of alternatives, *preference logics* aim at *reasoning about* preferences. A preference logic consists of a semantics and/or a formal system for interpreting relative preferences between logical formulas, or monadic, absolute preferences over formulas. The starting point of preference logics is that individuals often express relative or absolute preferences that refer not to isolated alternatives, but to logical formulas representing sets of alternatives. The central component of a preference logic is the *lifting* operator inducing preferences between formulas from preferences over single alternatives. At least two families of preference logics have been developed:

Logics of *ceteris paribus* Preferences. When an agent expresses a preference statement such as "I prefer to spend my summer holiday in Kentucky than in California", they surely do not mean that they prefer *any* summer holiday in Kentucky to *any* summer holiday in California; the preference statement does not preclude that they would prefer a sunny holiday in California to a rainy one in Kentucky. The principle at work when interpreting such preference statements is that the alternatives should be compared *all other things being equal (ceteris paribus)*, or more generally, all irrelevant properties being equal. A few milestones in *ceteris paribus* preference logics are [7–10]. Note that [10] also compares and attempts to reunify compact representation languages and preference logics.

Defeasible Preferences and Conditional Preference Logics. Consider the following statements: (1) I'd like to spend my weekend in Lexington; (2) if there is a storm warning on Lexington next weekend, then I prefer not to go. Statement (1) corresponds to a *defeasible*, *default* preference: it applies not only if we know that there is no risk of storm but more generally if there is no specific information about the weather forecast. This corresponds to assuming that the state of the world is *normal* (no storm warning); upon receiving the storm forecast, (1) is overridden by the more specific statement (2). Defeasible preferences fit the intuition as well as the natural language expression of preferences, and allow for their succinct and modular representation: succinct because they avoid to specify explicitly all the exceptional conditions in which a preference statement does not apply; modular because a set of such preference statements can be completed at any time, without generating an inconsistency — coming back to the latter example, if to (1) and (2) we later add (3) *if there is an exciting joint conference on algorithmic decision theory and logic programming in Lexington then I want to be there (independently of the weather)*, it will have higher priority than (2) in the 'doubly exceptional' circumstance "stormy weather and ADT-LPNMR". A neat way of formalizing these defeasible, conditional preferences consists in using *conditional logics*; some key papers on conditional preferences are [11–13].

3 Representing and Reasoning with Beliefs

Logic (and in particular, doxastic and epistemic logics) allow for distinguishing between objective facts and subjective beliefs. As there are numerous classes of problems in algorithmic decision theory where the decision maker(s) do not have a complete knowledge of the situation, logic has definitely a role to play here.

3.1 The External Perspective: Incomplete Knowledge of Agents' Preferences

Let us consider the point of view of an external agent that has to make decisions or to make predictions based on an incomplete, partial view of the agents' preferences: for instance, a recommender system in decision aid, a central authority in group decision making, or the modeller in game theory.

For the sake of brevity, let us focus on social choice. Often, the central authority in charge of computing the outcome (the winner of an election, an allocation of resources, etc.) has an incomplete knowledge of the agents' preferences, perhaps because the elicitation process was not conducted until its end, or because the voters could not report complete preferences. The central authority sees a set of possible worlds, each corresponding to a complete preference profile; an alternative is a possible (resp. necessary) winner if it is a winner in some (resp. all) possible worlds(s) [14].[2] Similar notions have been studied in fair division. It is clear that these notions originate in epistemic logic.

3.2 The Internal Perspective: Beliefs and Strategic Behaviour

A crucial issue in distributed multiagent systems is the impact of strategic behaviour on the 'social quality' of the reached state. Focusing on social choice, a tremendously high number of papers examine the conditions under which a mechanism that takes as input the agents' declared preferences can be manipulated by them, the computational complexity of finding a manipulation, and the impact of manipulation on social welfare. The assumption typically made is that agents have complete knowledge of the others' preferences. What if they have complex mutual beliefs, weaker than common knowledge, but stronger than zero

[2] See [15] for a review of existing work along this line.

knowledge? First steps towards handling such mutual beliefs in social choice have been made [16–19], but they remain preliminary. On the other hand, on reasoning about mutual beliefs in game theory there is an abundant literature, and even a series of workshops (*Logical Foundations of Decision and Game Theory*).

4 Logic for Problem Solving

4.1 Logical Encoding and Resolution of Decision Problems

Probably the widest use of logic in decision making contexts takes place in sequential decision making settings, or more generally in contexts where one has to search in a combinatorial space of solutions. The paradigmatic example is *planning as satisfiability* [20]: the planning problem (initial state, action effects, goal, horizon) is translated into a set of propositional clauses, which is fed to a SAT solver; the model found by the solver (if any) is translated back to a plan. The framework has been extended to planning with nondeterministic actions.[3] Answer set programming is also a natural and efficient tool for expressing and solving planning problems [22,23] and for multicriteria optimisation [24].

4.2 Automated Theorem Proving and Discovery

Open research questions can be addressed using computer-aided theorem proving techniques. The role of computer science here is not to help solving a decision making problem, but to (re)prove theorems and/or discover new ones. Automated theorem proving and discovery is especially helpful in branches of decision making where combinatorial structures prevail, such as decision theory over discrete domains, social choice theory, cooperative or noncooperative game theory. Some examples are an automated proof of Arrow's theorem [25,26], impossibility theorems about pure Nash equilibria in two-person games [27], about ranking sets of alternatives [28], or about strategyproofness and participation in voting [29,30]. Also, the modelling of social choice mechanisms in modal logic [31,32] is related to this research line.

References

1. Lang, J., Xia, L.: Voting in combinatorial domains. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice. Cambridge University Press (2015, in Press)
2. Lang, J.: Logical representation of preferences. In: Bouyssou, D., Dubois, D., Prade, H., Pirlot, M. (eds.) Decision Making Process: Concepts and Methods, pp. 321–364. Wiley-ISTE, London (2009)
3. Delgrande, J.P., Schaub, T., Tompits, H., Wang, K.: A classification and survey of preference handling approaches in nonmonotonic reasoning. Comput. Intell. **20**, 308–334 (2004)

[3] See [21] for a survey.

4. Brewka, G.: Preferences in answer set programming. In: Marín, R., Onaindía, E., Bugarín, A., Santos, J. (eds.) CAEPIA 2005. LNCS (LNAI), vol. 4177, pp. 1–10. Springer, Heidelberg (2006)
5. Brewka, G., Niemelä, I., Truszczynski, M.: Preferences and nonmonotonic reasoning. AI Mag. **29**, 69–78 (2008)
6. Brewka, G., Delgrande, J.P., Romero, J., Schaub, T.: asprin: customizing answer set preferences without a headache. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, Austin, Texas, USA, 25–30 January 2015, pp. 1467–1474 (2015)
7. von Wright, G.: The Logic of Preference. Edinburgh University Press, Edinburgh (1963)
8. Hansson, S.O.: The Structure of Values and Norms. Cambridge University Press, Cambridge (2001)
9. van Benthem, J., Girard, P., Roy, O.: Everything else being equal: a modal logic for Ceteris Paribus preferences. J. Philos. Logic **38**, 83–125 (2009)
10. Bienvenu, M., Lang, J., Wilson, N.: From preference logics to preference languages, and back. In: Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, 9–13 May 2010
11. Boutilier, C.: Toward a logic for qualitative decision theory. In: Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR 1994), Bonn, Germany, 24–27 May 1994, pp. 75–86 (1994)
12. Lang, J.: Conditional desires and utilities: an alternative logical approach to qualitative decision theory. In: Proceedings of the 12th European Conference on Artificial Intelligence, Budapest, Hungary, 11–16 August 1996, pp. 318–322 (1996)
13. Lang, J., van der Torre, L.W.N., Weydert, E.: Hidden uncertainty in the logical representation of desires. In: IJCAI 2003, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, 9–15 August 2003, pp. 685–690 (2003)
14. Konczak, K., Lang, J.: Voting procedures with incomplete preferences. In: Multidisciplinary Workshop on Advances in Preference Handling, pp. 124–129 (2005)
15. Boutilier, C., Rosenschein, J.: Incomplete information and communication in voting. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice. Cambridge University Press (2015, in Press)
16. Chopra, S., Pacuit, E., Parikh, R.: Knowledge-theoretic properties of strategic voting. In: Alferes, J.J., Leite, J. (eds.) JELIA 2004. LNCS (LNAI), vol. 3229, pp. 18–30. Springer, Heidelberg (2004)
17. van Ditmarsch, H., Lang, J., Saffidine, A.: Strategic voting and the logic of knowledge. In: Proceedings of 14th TARK, ACM (2013)
18. Meir, R., Lev, O., Rosenschein, J.: A local-dominance theory of voting equilibria. In: ACM Conference on Economics and Computation, EC 2014, Stanford, CA, USA, 8–12 June 2014, pp. 313–330 (2014)
19. Conitzer, V., Walsh, T., Xia, L.: Dominating manipulations in voting with partial information. In: Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, 7–11 August 2011
20. Kautz, H., Selman, B.: Planning as satisfiability. In: ECAI, pp. 359–363 (1992)
21. Rintanen, J.: Planning and sat. In: Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.) Handbook of Satisfiability, pp. 483–504. IOS Press, Amsterdam (2009)
22. Lifschitz, V.: Answer set programming and plan generation. Artif. Intell. **138**, 39–54 (2002)

23. Eiter, T., Faber, W., Leone, N., Pfeifer, G., Polleres, A.: A logic programming approach to knowledge-state planning, II: the $\mathrm{DLV}^{\mathrm{k}}$ system. Artif. Intell. **144**, 157–211 (2003)
24. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Multi-criteria optimization in answer set programming. In: Technical Communications of the 27th International Conference on Logic Programming, p. 1 (2011)
25. Nipkow, T.: Social choice theory in HOL: Arrow and Gibbard-Satterthwaite. J. Autom. Reasoning **43**, 289–304 (2009)
26. Tang, P., Lin, F.: Computer-aided proofs of Arrow's and other impossibility theorems. Artif. Intell. **173**, 1041–1053 (2009)
27. Tang, P., Lin, F.: Discovering theorems in game theory: two-person games with unique pure nash equilibrium payoffs. Artif. Intell. **175**, 2010–2020 (2011)
28. Geist, C., Endriss, U.: Automated search for impossibility theorems in social choice theory: ranking sets of objects. J. Artif. Intell. Res. **40**, 143–174 (2011)
29. Brandt, F., Geist, C.: Finding strategyproof social choice functions via SAT solving. In: AAMAS 2014, pp. 1193–1200. IFAAMAS (2014)
30. Brandl, F., Brandt, F., Geist, C., Hofbauer, J.: Strategic abstention based on preference extensions: positive results and computer-generated impossibilities. In: IJCAI 2015 (2015)
31. Ågotnes, T., van der Hoek, W., Wooldridge, M.: On the logic of preference and judgment aggregation. Auton. Agents Multi-Agent Syst. **22**, 4–30 (2011)
32. Troquard, N., van der Hoek, W., Wooldridge, M.: Reasoning about social choice functions. J. Philos. Logic **40**, 473–498 (2011)

http://www.springer.com/978-3-319-23263-8

Logic Programming and Nonmonotonic Reasoning
13th International Conference, LPNMR 2015, Lexington,
KY, USA, September 27-30, 2015. Proceedings
Calimeri, F.; Ianni, G.; Truszczynski, M. (Eds.)
2015, XIV, 574 p. 76 illus., Softcover
ISBN: 978-3-319-23263-8