Università di Padova – Dipartimento di Matematica Pura ed Applicata

Scuole di Dottorato in Matematica Pura e Matematica Computazionale

Seminario Dottorato 2006/07

Presentazione

Questo documento è il resoconto dell'attività del Seminario Dottorato 2006/07. Svolto in forma regolare per la prima volta durante quest'anno accademico, il Seminario Dottorato fornisce un'opportunità a dottorandi, assegnisti e borsisti in Matematica Pura e Computazionale di comunicare le proprie ricerche in modo interessante per un pubblico matematicamente ben istruito ma non specialista, traendone giovamento in termini di autocoscienza di ciò che si sta studiando. La cura con la quale i seminari sono stati preparati (di cui è testimonianza anche la redazione di queste note riassuntive), e la generalmente buona partecipazione dei compagni di corso e di altri colleghi di volta in volta interessati, ci autorizzano ad archiviare in modo soddisfacente questo ciclo, e a programmarne un proseguimento anche per l'avvenire.
Un grazie va a tutti coloro che hanno contribuito alla riuscita dell'iniziativa, in primo luogo naturalmente a chi ha accettato di svolgere i seminari e di stendere queste note.

Padova, 2 luglio 2007

Corrado Marastoni, Tiziano Vargiolu

Problems with Preferences and Uncertainty

Maria Silvia Pini (*)

Abstract. Many real-life problems are characterized by preferences and uncertain events, which cannot be controlled by the agent. We present formalisms that model problems with various kinds of preferences and with uncertainty, we study the properties of such formalism, and we consider scenarios where several agents express simultaneously their preferences. [Keywords: preferences, uncertainty, multi-agent systems.]

Sunto. *Problemi con preferenze ed incertezza.* Molti problemi reali sono caratterizzati da preferenze e da eventi incerti che non possono essere controllati dall'utente. Presentiamo dei formalismi che modellano problemi con vari tipi di preferenze e l'incertezza, studiamo le proprietà di questo formalismi e consideriamo scenari in cui più utenti esprimono contemporaneamente le loro preferenze.

1 Motivation

Preferences are ubiquitous in real life. In fact, most problems are over-constrained and would not be solvable if we insist that all their requirements are strictly met, hence it is more reasonable to express their requirements in a soft way, i.e., via preferences. Moreover, many problems are more naturally described via preferences rather than hard statements.

Preferences come in many kinds. In some cases it could be natural to express preferences in quantitative terms, while in other situations it could be better to use qualitative statements. Moreover, preferences can be unconditional or conditional. Finally, preferences can model priorities, rankings, different levels of importance, desires or rejection levels.

Preferences can help whenever the task involves decision making and/or knowledge representation. They are essential to treat reasoning about action and time, planning diagnosis and configuration. Preferences are the key to understand the non-crisp aspect of many human behaviors. For example, in mathematical decision theory, preferences (often expressed as utilities) are used to model people's economic behavior. In Artificial Intelligence (AI), preferences help to capture agents' goals. In philosophy, preferences are used to reason about values, desires, and duties. Thus, the representation and handling

(*)Dottorato in Matematica Computazionale, XIX ciclo. Seminario tenuto il 28 marzo 2007.
e-mail: `mpini@math.unipd.it` – Address: University of Padova, Department of Pure and Applied Mathematics, Via Trieste, 63 - 35121 Padua

of preferences should be available and efficient in any sophisticated automated reasoning tool. Preferences are gaining more and more attention in AI, in particular in the Constraint Programming (CP) area also in connection with Operations Research (OR). AI permits complex preference representations and thus allows to reason with and about preferences, providing a new perspective for formalizing preference information in qualitative and quantitative way, that is essential for many decision making problems.

Preferences are one type of soft information present in real-life problems. Another important feature, which arises in many real world problems, is uncertainty. In fact, many problems are characterized by uncertain parameters which are not under the user's direct control, but that can be decided only by Nature. An example of an uncertain parameter is, in the context of satellite scheduling or weather prediction, the time when clouds will disappear, which can be decided only by Nature. Another example in which uncertainty occurs is a scheduling problem, which constrains the order of execution of various activities, where the duration of some activity is uncertain. In this case the goal is to define a schedule which is the most robust with respect to uncertainty.

Uncertainty can be represented in several ways. In some problems the user can be completely ignorant about the occurrence of the uncertain events, in others he can have additional information, which can be more or less precise regarding the occurrence of uncertain events [4,5].

Preferences and uncertainty often coexist in real-word problems. Consider for example a scheduling problem with uncertain durations, which is over-constrained. It would be impossible to solve it if we insist that all its requirements are strictly met. Therefore, it is more reasonable to express (at least some of) its requirements as preferences rather than hard statements. Doing so, we obtain a problem defined by preferences and uncertainty, where the solutions are schedules with different levels of desirability. The goal is then to find solutions with the highest level of desirability which are also robust with respect to uncertain durations.

Since preferences and uncertainty are very often the core of real-life problems, it is important to model faithfully these two aspects, both for problems involving a single agent and for problems regarding multiple agents. While there are several formalisms to handle some notions of preferences and/or uncertainty, much work still needs to be done to handle them in a general and efficient way. We give a contribution in this direction.

2 Objectives

There are many issues to be addressed about preferences and uncertainty. The main issue is preference and uncertainty specification and representation, i.e., which formalisms can be used to model the preferences of an agent and the uncertainty of the problem. In this respect, contributions have been brought from studying the axiomatic properties of preferences, as well as logics of preferences or their topological and algebraic structures [1] and from defining formalisms for representing various kinds of uncertainty [5]. In a multi-agent scenario, instead, core issues are preference composition, merging and aggregation, as well as preference elicitation, conflict resolution and belief revision.

Our goal is to define and study formalisms that can model problems with many kinds of preferences and/or uncertainty, to study properties of such formalisms, and to develop tools to solve such problems. Moreover, we want to deal also with scenarios where preferences are expressed by several agents and where preference aggregation is therefore needed to find the optimal solutions.

In order to achieve this objective, we start by defining formalisms for expressing preferences of a single agent in presence of uncertainty. We start considering problems where preferences are expressed in a quantitative non conditional way and where uncertainty is characterized by lack of data or imprecise knowledge. In some formalisms for dealing with preferences and uncertainty, uncertainty is expressed in terms of probability theory [5]. We consider a different form of uncertainty, less precise than the probabilistic one, since we intend to model scenarios where probabilistic estimates are not available. We define a formalism for handling preferences and this kind of uncertainty, and algorithms for solving them. To achieve this goal, we exploit two formalisms: the semiring-based soft constraint formalism [1] to deal with preferences, and possibility theory [4] to reason with uncertainty.

Generally speaking, a soft constraint is just a classical constraint plus a way to associate, either to the entire constraint or to each assignment of variables, a certain element, which is usually interpreted as a level of preference or importance. Such levels are usually ordered and the order reflects the idea that some levels are better than others. Moreover, one has also to say, via suitable combination operators, how to obtain the level of preference of a global solution from the preferences in the constraint.

Many formalisms have been developed to describe one or more classes of soft constraints. For instance, consider fuzzy CSPs, where the crisp constraints are extended with a level of preference represented by a real number between 0 and 1, or probabilistic CSPs, where the probability to be in the real problem is assigned to each constraint. We choose to use one of the most general frameworks to deal with soft constraints [1]. The framework is based on a semiring structure that is equipped with the operations needed to combine the constraints present in the problem and to choose the best solutions. According to the choice of the semiring, this framework is able to model all the specific soft constraint notions mentioned above. The semiring-based soft constraint framework provides a structure capable of representing in a compact way problems with preferences.

For handling uncertainty we consider possibility theory, which is a mathematical theory for dealing with a certain type of uncertainty. This theory is an alternative to probability theory. It can be seen as an imprecise probability theory. Possibility theory has been introduced as an extension of the theory of fuzzy sets and fuzzy logic and many contributions to its development have been presented, for example, in [4].

Another issue that we consider is the representation of bipolarity. Bipolarity is an important focus of research in several domains, e.g. psychology, multi-criteria decision making, and more recently in AI (argumentation and qualitative reasoning). Preferences on a set of possible choices are often expressed in two forms: positive and negative statements. In fact, in many real-life situations agents express what they like and what they dislike, thus often preferences are bipolar. Starting from this observation, we define a formalism for handling quantitative (unconditional) preferences, which is able to represent positive and negative statements, and also to deal with uncertainty. Starting from an existing formalism for handling negative preferences, i.e., soft constraints [1], we extend it to handle also positive preferences. The aim is to handle bipolar preferences in a way which is as similar as possible to what naturally happens in real-life scenarios. That is, combining two negative statements should be even worse, combining two positive statements should be even better, and combining a positive with a negative statement should be positive if the positive statement is stronger than the negative one, and negative otherwise. Moreover, we want to able to express indifference, i.e., a preference which is neither positive nor negative.

In many situations, we need to represent and reason about simultaneous preferences of several agents. To aggregate agents' preferences, which in general express a partial order over the possible outcomes, we can query each agent in turn and collect together the results. Hence, we can see preference aggregation in terms of voting, which is a topic widely studied in Operations Research. In this context, we study classical properties such as fairness and non-manipulability, and we consider classical results on fairness of social [6] welfare functions as Arrow's impossibility theorem and results on non-manipulability of social choice functions as Gibbard-Satterthwaite's theorem [6]. The main difference is that, in contrast to what is assumed in social welfare scenarios, our agents describe their preference using partial orders and not total orders, i.e., they can consider incomparable pairs of outcomes, which are too dissimilar to be compared. We study if results similar to the ones of social choice and social welfare settings still hold, by suitably adapting some of their assumptions to deal with incomparability.

Finally, we consider uncertainty in a multi-agent scenario. We consider a multi-agent setting where agents can hide some of their preferences. In a preference ordering, the relationship between some pairs of outcomes may not be specified. For example, agents may have privacy concerns about revealing their complete preference ordering or, as in the context of preference elicitation, preferences have not been fully elicited [3]. In this context it is interesting to determine the complexity of computing the outcomes which are always optimal, or optimal in at least one way in which incompleteness is resolved. Moreover, if this computation is difficult, it is useful to find cases where this computation is easy. Regarding this topic, we investigate such complexity results both in general and for specific preference aggregation systems, and we analyze the issue of manipulation in this context.

3 Results

We have followed the research lines outlined in the previous section, and we have obtained the following main results.

We have started by considering a special case of quantitative preferences, i.e., fuzzy preferences, and we have considered an existing technique to integrate such preferences with uncertainty, which uses possibility theory [4]. We have shown that the integration provided by this technique is too tight since the resulting ordering over solutions does not allow one to discriminate between solutions which are highly preferred but assume unlikely events and solutions which are not preferred but robust with respect to uncertainty. Thus, while following the same basic idea of translating uncertainty into fuzzy constraints, we have proposed an algorithm which allows us to observe separately the preference and the robustness of the solutions. Moreover, we have defined suitable semantics for ordering the solutions in a more or less risky way with respect to uncertainty. Then, for finding optimal solutions according to the different semantics, we have developed a solver which exploits branch and bound techniques [8]. Moreover, we have defined a more general formalism for handling different kinds of quantitative preference, proving that some desirable properties continue to hold [9]. This has allowed us to handle the coexistence of preferences and uncertainty in a more general setting.

We have also defined a formalism to handle positive and negative preferences, which reflects the natural behaviour that the combination of positive and negative statements has in real-life scenarios. For doing so, we have first shown that the negative preferences are handled by the semiring-based formalism of soft constraints. Then, we have introduced a new algebraic structure for handling positive preferences, which has properties similar to semirings. Hence, we have defined a new mathematical structure for handling both positive and negative preferences by linking the positive and the negative structures in a suitable way, so that combination of positive preferences produces a better positive preference, the combination of negative preferences produces a worse negative preference, and the combination of positive and negative preferences produces a preference which is better than or equal to the negative preference and worse than or equal to the positive one. We have studied the properties of this formalism and we have defined a solver to solve such problems [2]. Moreover, we have generalized this solver to handle also uncertainty.

We have then considered scenarios where several agents express their preferences via a partial order over the possible outcomes. We have seen each agent as voting if an outcome dominates another one. Thus, we have considered preference aggregation in terms of voting, analyzing some of the main results concerning fairness and non-manipulability [6]. We have shown that they can be generalized to preference aggregation systems, where agents can express also incomparability between pair of outcomes [10,11].

We have finally considered scenarios where agents, for example for privacy reasons, decide to hide some of their preferences. We have determined the computational complexity of computing optimal outcomes, where optimality has the meaning of being always the best outcome (regardless of how incompleteness is resolved), or in at least one possible complete world. We have shown that computing such outcomes is in general difficult, and we have determined cases where such a problem is tractable. Moreover, we have shown how the computation of such outcomes can be useful for deciding when preference elicitation is over, which is in general a difficult problem [12]. Finally, we have investigated other tractability and intractability results for a specific voting rule, i.e, the sequential majority voting. Such a rule performs a sequence of pairwise comparisons between two candidates along a binary tree, and the winner depends on the chosen sequence. We have focused on candidates that will win in some sequences or in all sequences and we have shown that in general it is easy to find them, while it is difficult if we insist that the tree is balanced. We have interpreted this difficulty in terms of difficulty for the chair to manipulate [7].

Acknowledgments

I would like to thank my coauthors, Francesca Rossi, Kristen Brent Venable, Toby Walsh, Stefano Bistarelli and Jerome Lang.

References

[1] S. Bistarelli, U. Montanari, and F. Rossi, *Semiring-based constraint solving and optimization.* Journal of the ACM, 44(2):201236, mar 1997.

[2] S. Bistarelli, M. S. Pini, F. Rossi, and K. B. Venable, *Bipolar preference problems.* In Proceedings of ECAI 2006, 17th European Conference on Artificial Intelligence, pages 705706. IOS Press, 2006.

[3] V. Conitzer and T. Sandholm, *Vote elicitation: complexity and strategy- proofness.* In Proceedings of AAAI-02, pages 392397, 2002.

[4] D. Dubois, H. Fargier, and H. Prade, *Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty.* Appl. Intel l., 6(4):287309, 1996.

[5] H. Fargier and J. Lang, *Uncertainty in constraint satisfaction problems: a probabilistic approach.* In M. Clarke, R. Kruse, and S. Moral, editors, ECSQARU93, volume 747 of LNCS, pages 97104. Springer, 1993.

[6] J. S. Kelly, "Social Choice Theory: An Introduction". Springer Verlag, Berlin, 1987.

[7] J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh, *Winner determination in sequential majority voting.* In IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, pages 13721377. AAAI Press, 2007.

[8] M. S. Pini, F. Rossi, and K. B. Venable, *Possibility theory for reasoning about uncertain soft constraints.* In Proceedings of the 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU 2005, volume 3571 of LNCS. Springer, 2005.

[9] M. S. Pini, F. Rossi, and K. B. Venable, *Uncertainty in soft constraint problems.* In Proceedings of the International Conference on Intelligent Agents, Web Technologies and Internet Commerce (IAWTIC 2005), pages 583589. IEEE Computer Society, 2005.

[10] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh, *Aggregating partially ordered preferences: impossibility and possibility results.* In Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2005), pages 193206. National University of Singapore, 2005.

[11] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh, *Strategic voting when aggregating partially ordered preferences.* In Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2006, pages 685687. ACM Press, 2006.

[12] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh, *Incompleteness and incomparability in preference aggregation.* In IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, pages 14641469. AAAI Press, 2007.