Graphical Representation of Ordinal Preferences: Languages and Applications

Jérôme Lang

LAMSADE, Université Paris-Dauphine, 75775 Paris Cedex 16, France
lang@lamsade.dauphine.fr

1 Introduction

The specification of a decision making problem includes the agent's preferences on the available alternatives. The choice of a *model* of preferences (*e.g.*, utility functions or binary relations) does not say how preferences should be *represented* (or *specified*). A naive idea would consist in writing them *explicitly*, simply by enumerating all possible alternatives together with their utility (in the case of cardinal preferences) or the list of all pairs of alternatives contained in the relation (in the case of ordinal preferences). This is feasible in practice only when the number of alternatives is small enough with respect to the available computational resources. This assumption is often unrealistic, in particular when the set of alternatives has a combinatorial (or multiattribute) structure, *i.e.*, when each alternative consists of a tuple of values, one for each of a given set of *variables* (or *attributes*): in this case, the set of alternatives is the Cartesian product of the value domains, and its cardinality grows exponentially with the number of variables.

For such combinatorial domains, we need a language allowing to express preferences as succinctly as possible. Such *compact preference representation languages* have been particularly studied in the Artificial Intelligence research community. A significant number of these languages are called "graphical', because they consist of a graphical component describing preferential dependencies between variables, together with a collection of local preferences on single variables or small subsets of variables, compatible with the dependence structure.

In this short paper we will focus on graphical languages for *ordinal preferences*, and especially on CP-nets and their extensions and variants[1]. After giving an brief presentation of this family of languages, we will show how they can be used for individual, collective or distributed decision making.

2 Graphical Languages for Ordinal Preference Representation

CP-Nets

Let $V = \{X_1, \ldots, X_n\}$ be a set of *variables*, also called *attributes*, with their associated finite *domains* $D_1, \ldots, D_n$. By $\mathcal{V} = \times_{X_i \in V} D_i$, we denote the set of all complete assignments, called *outcomes* (or *alternatives*). For $X \subseteq V$, we let $\mathcal{X} = \times_{X_i \in X} D_i$. For any

[1] However, at some places in the text we will refer to a graphical languages for *cardinal* preference representation such as *GAI*-nets [2], which specify local utility functions on small subsets of variables.

M. Croitoru, S. Ferré, and D. Lukose (Eds.): ICCS 2010, LNAI 6208, pp. 3–9, 2010.

disjoint subsets X and Y of V, the concatenation of assignments $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, denoted xy, is the $(X \cup Y)$-assignment which assigns to variables in X (resp. Y) the value assigned by x (resp. y).

A *(strict) preference relation* $\succ$ is an irreflexive and transitive (thus asymmetric) binary relation over $\mathcal{V}$. A strict preference relation $\succ$ is *linear* if it is connected, that is, for every $x \neq y$ we have either $x \succ y$ or $y \succ x$.

Let $\{X,Y,Z\}$ be a partition of the set of variables and $\succ$ a strict preference relation. X is *preferentially independent of* Y *given* Z (with respect to $\succ$) if and only if for all $x_1,x_2 \in \mathcal{X}$, $y_1,y_2 \in \mathcal{Y}$, $z \in D_Z$, we have $x_1y_1z \succ x_2y_1z$ if and only if $x_1y_2z \succ x_2y_2z$ [15].

A CP-net [5] is composed of a directed graph representing the preferential dependencies between variables, and a set of conditional preference tables expressing, for each variable, the local preference on the values of its domain given the possible combination of values of its parents. Formally, a *CP-net* over $V = \{X_1,\dots,X_n\}$ is a pair $\mathcal{N} = \langle G,P \rangle$ where G is a directed graph over $X_1,\dots,X_n$ and P is a set of conditional preference tables $CPT(X_i)$ for each $X_i \in V$. For variable X_i, we denote by $Par(X_i)$ the set of parents of X_i in G and we let $NonPar(X_i) = \mathcal{V} \setminus (\{X_i\} \cup Par(X_i))$. Each conditional preference table associates a linear order on D_i with each instanciation u of $Par(X_i)$, denoted by $u :>$. The edges of $\mathcal{G}$ express preferential independencies: every variable is preferentially independent of its non-parents in G given its parents.

The semantics of a CP-net is defined as follows. A strict preference relation $\succ$ satisfies $\mathcal{N}$ if for all variables X_i, values $x_i,x_i' \in D_i$, assignments u of $Par(X_i)$, and assignments z of $NonPar(X_i)$, we have $ux_iz \succ ux_i'z$ if and only if $CPT(X_i)$ contains $u : x_i > x_i'$. A CP-net is *satisfiable* if there is some preference relation satisfying it. For any satisfiable CP-net $\mathcal{N}$, $\succ_{\mathcal{N}}$ is the smallest preference relation that satisfies $\mathcal{N}$.

Example 1. A user is looking for an airplane ticket. There are three variables: T (time of the flight), with possible values d (day) and n (night); S (stopover), with possible values s (yes) and $\bar{s}$ (no); and C (company), with two possible values c_1 and c_2. The agent has the following preferences: (a) she prefers a day flight to a night flight, unconditionally; (b) for a day flight she prefers to make a stopover; for a night flight she'd rather prefer not; (c) for a day flight with a stopover she prefers company c_1 because it implies spending a few hours in an airport she likes; in all other cases she prefers c_2.

The user's preferences are expressed by the CP-net $\mathcal{N}$ depicted below, together with the induced preference relation $\succ_{\mathcal{N}}$.

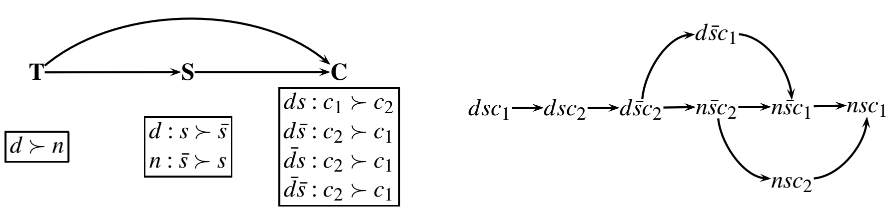

Many works on CP-nets make the additional assumption that the graph $\mathcal{G}$ is *acyclic*. Under this assumption, the CP-net is satisfiable, and the associated requests, consisting in comparing two outcomes or in searching for a non-dominated outcome, are computationally reasonable [6].

The preference relation $\succ_{\mathcal{N}}$ induced from a CP-net $\mathcal{N}$ is generally not complete. The complete preference relations extending $\succ_{\mathcal{N}}$ can be viewed as possible models of the user's preferences, and any preference assertion that holds in all such models can be viewed as a consequence of the CP-net [6]. Finally, for any complete preference relation $\succ$ there exists a satisfiable CP-net $\mathcal{N}$ such that $\succ$ extends $\succ_{\mathcal{N}}$ (note however that the dependence graph of $\mathcal{N}$ may contain cycles).

Extensions and Variants of CP-Nets

TCP-nets [9] enrich CP-nets by allowing the expression of relative importance statements between variables. CP-theories [20] are still more general; they allow conditional preference statements on the values of a variable, together with a set of variables that are allowed to vary when interpreting the preference statement. The language considered in [21] is even more general: there the preference statements do not compare single values of variables but tuples of values of different variables. *Conditional Importance Networks* [7] express monotonic preferences between sets of goods, *ceteris paribus*.

3 Individual Decision Making

Quoting from [6], the main purpose of any preference representation language is to support answering various queries about the preferences of the decision maker. Two fundamental types of queries are (a) *outcome comparison* (given two outcomes, determine if one of them dominates the other) and (b) *outcome optimization* (determine the best outcome(s)). Now, in many decision making problems, not all outcomes are feasible. A *constrained CP-net* consists of a CP-net $\mathcal{N}$ and a set of constraints Γ. Constraints can be expressed compactly using some representation language, typically CSPs (or, in the case of Boolean variables, propoisitional logic). Any outcome satisfying Γ is said to be *feasible*. The goal is to find an outcome α that is both feasible and undominated, *i.e.*, such that there is no feasible outcome β such that $\beta \succ_{\mathcal{N}} \alpha$.

Consider Example 1 again, and let us add the constraint that no day flight with a stopover is available: $T = d \Rightarrow S = \bar{s}$, and the constraint that company c_2 has only night flights: $C = c_2 \Rightarrow T = n$. The outcome dsc_1, which was the optimal outcome of $\succ_{\mathcal{N}}$, is now unfeasible. The new undominated outcomes are $d\bar{s}c_1$ and $n\bar{s}c_2$.

A different way of defining optimal solutions for constrained CSPs is suggested in [18]: α dominates β if there is a flipping sequence from α to β *that goes through feasible outcomes only*, and again we look for non-dominated feasible outcomes. Back to Example 1, suppose that we have the constraint $C = c_2 \Rightarrow T = n$. According to [18], dsc_1 and $n\bar{s}c_2$ are undominated, whereas only dsc_1 is undominated according to [6].

Constrained optimization is particularly relevant for configuration problems (see for instance [11] for an application to personalized configuration of web-page content). Another form of constrained optimization can come from the fact that an outcome is feasible only if there is a plan that realizes it; in [8], preferences among goals states are specified using a TCP-net, and one looks for a plan resulting in an optimal outcome, that is, an state α such that no other reachable state dominates α. A last example of using CP-nets for individual decision making is [4], which describes an information retrieval approach where CP-nets are used for expressing preferences over documents.

4 Group Decision Making

An important problem in computational social choice is voting on a combinatorial set of alternatives: a number of 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. 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. As soon as voters have preferential dependencies between issues, it is generally a bad idea to decompose the problem into a set of smaller problems, each one bearing on a single issue: doing so may lead to *multiple election paradoxes* [10].

The input of a voting problem is typically a collection of linear preference relations. Therefore, graphical language for ordinal preferences, which express preferences locally and exploit preferential independencies, are particularly well-suited to the design of methods for collective decision making on combinatorial domains. There are two different ways of making use of graphical languages in such contexts.

The first way consists in *eliciting preferences globally, and then aggregating them*. For instance, if we are using CP-nets, we first elicit a CP-net for each voter, these CP-nets are then aggregated so as to output an outcome or a set of outcomes. See [19] and [22] for two approaches for CP-net aggregation (see also [12] for group decision making based on the aggregation of GAI-nets).

The second way consists in proceeding *sequentially*: at every step, we elicit the voters' preferences about a single variable, we use a local rule to compute the value chosen for this variable, and this value is communicated to all voters. For this we need to make an important domain restriction, namely, that there exists an order on variables, say $X_1 > \ldots > X_p$, such that for every voter and for every i, X_i is preferentially independent of $X_{i+1},\ldots,X_p$ given $X_1,\ldots,X_{i-1}$ [17].

Consider the following example. We have a variable F(ood), with three possible values m(eat), f(ish) and v(egetarian) and a variable W(ine), with two possible values w(hite) and r(ed). Assume we have seven voters, all with unconditional preferences on F (and possible preferences on W depending on the value of F) whose conditional preferences are as follows:

1 voter: $m \succ f \succ v$; $m : r \succ w$; $f,v : w \succ r$
2 voters: $m \succ v \succ f$; $r \succ w$
2 voters: $f \succ v \succ m$; $m : r \succ w$; $f,v : w \succ r$
2 voters: $v \succ f \succ m$; $w \succ r$

If we apply the second method, assume that for F we use the Borda rule, which assigns a score or +2 (resp. +1, 0) to a candidate every time it is ranked first (resp. second, last) by a voter, and chooses the candidate that maximizes the sum of scores obtained for the different voters. Then preferences about food are elicited, the chosen value is v, then the voters' preferences about wine given that $F = v$ are elicited, and finally, after using the majority rule, we get that the collectively chosen value of W is w. If we apply the first method, we first aggregate the CP-nets, which, using the same local rules, results in the collective CP-net $v \succ f \succ m$; $m : r \succ w$; $f,v : w \succ r$, whose unique undominated outcome is $F = v, W = w$, which again is the resulting collective decision (note that there are cases where the results given by both methods would have been different). Whereas the second method is cheaper in communication and computation, its range of applicability is much more restricted. Indeed, if the first voter's conditional preferences were $r : m \succ f \succ v$; $w : f \succ v \succ m$, $m : r \succ w$; $f, v : w \succ r$, then asking her to express her preferences on F would make her feel ill at ease, since her preferences on F depend on the value of W.

5 Game Theory

Game theory attempts to analyze formally strategic interactions between agents. Roughly speaking, a static game consists of a set of agents, and for each agent, a set of possible strategies and a utility function mapping every possible combination of strategies to a real value. Utility functions are usually represented explicitly, by listing the values for each combination of strategies. However, the number of utility values which must be specified, that is, the number of possible combinations of strategies, is exponential in the number of players, which renders such an explicit way of representing the preferences of the players unreasonable when the number of players is more than a few units. This becomes even more problematic when the set of strategies available to an agent consists in assigning a value from a finite domain to each of a given set of variables. In these cases, compact representation are once again useful.

Now, several key notions in game theory, such as pure strategy Nash equilibria or dominated strategies, need only ordinal preferences. Graphical languages for ordinal preference representation have been used in some places for representing and analyzing games, such as Section 5 of [3]: there, a game is described as a set of agents, a set of (binary) variables, a control assignment function assigning each variable to an agent, and finally, a compact description of the agents' preferences under the form of a collection of CP-nets. The structure of the agent's preferences can sometimes guarantee the existence or the unicity of a pure Nash equilibrium. For instance, assume we have two agents $\{1,2\}$ and two variables A and B, with domains $\{a, \bar{a}\}$ and $\{b, \bar{b}\}$, such that 1 controls A and 2 controls B. Preferential dependencies are as follows:

	situation 1	situation 2	situation 3
agent 1	$A \longrightarrow B$	$A \longrightarrow B$	$A \longleftarrow B$
agent 2	$A \longleftarrow B$	$A \longrightarrow B$	$A \longrightarrow B$

In situation 1, whatever the local preferences of 1 and 2, there is a unique equilibrium consisting of 1's preferred value of A and 2's preferred value of B. In situation 2, again there is a unique equilibrium, consisting of 1's preferred value of A and 2's preferred value of B given 1's preferred value of A. In situation 3, neither the existence nor the unicity of a pure Nash equilibrium is guaranteed. For instance, if the local preferences of 1 are $\bar{b} \succ b$, $b : a \succ \bar{a}$, $\bar{b} : \bar{a} \succ a$, and the local preferences of 2 are $a \succ \bar{a}$, $a : b \succ \bar{b}$, $\bar{a} : \bar{b} \succ b$, then the game has two pure Nash equilibria, namely ab and $\bar{a}\bar{b}$ [3].

A different connection between CP-nets and games appears in [1], where CP-nets are viewed as games in normal form and vice versa: each player corresponds to a variable of the CP-net, whose domain is the set of actions available to the player.

Normal form games with *cardinal* preferences have received even more attention: graphical games [16,14,13] specify, for each agent, the set of all players that have an influence on her; the utility of player i is compactly represented by a utility table that specifies a value for each combination of actions of the players on which i depends.

References

1. Apt, K.R., Rossi, F., Venable, K.B.: CP-nets and Nash equilibria. In: Proceedings of Third International Conference on Computational Intelligence, Robotics and Autonomous Systems (CIRAS 2005). Elsevier, Amsterdam (2005)
2. Bacchus, F., Grove, A.: Graphical models for preference and utility. In: Proceedings of UAI 1995, pp. 3–10 (1995)
3. Bonzon, E., Lagasquie-Schiex, M.-C., Lang, J., Zanuttini, B.: Compact preference representation and Boolean games. Autonomous Agents and Multi-Agent Systems 1(18), 1–35 (2009)
4. Boubekeur, F., Boughanem, M., Tamine-Lechani, L.: Towards flexible information retrieval based on CP-nets. In: Larsen, H.L., Pasi, G., Ortiz-Arroyo, D., Andreasen, T., Christiansen, H. (eds.) FQAS 2006. LNCS (LNAI), vol. 4027, pp. 222–231. Springer, Heidelberg (2006)
5. Boutilier, C., Brafman, R., Domshlak, C., Hoos, H., Poole, D.: CP-nets: a tool for representing and reasoning with conditional ceteris paribus statements. JAIR 21, 135–191 (2004)
6. Boutilier, C., Brafman, R., Domshlak, C., Hoos, H., Poole, D.: Preference-based constrained optimization with CP-nets. Computational Intelligence 20(2), 137–157 (2004)
7. Bouveret, S., Endriss, U., Lang, J.: Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In: Proceedings of IJCAI 2009 (2009)
8. Brafman, R., Chernyavsky, Y.: Planning with goal preferences and constraints. In: Proceedings of ICAPS 2005, pp. 182–191 (2005)
9. Brafman, R., Domshlak, C., Shimony, S.E.: On graphical modeling of preference and importance. JAIR 25, 389–424 (2006)
10. Brams, S., Kilgour, D., Zwicker, W.: The paradox of multiple elections. Social Choice and Welfare 15(2), 211–236 (1998)
11. Domshlak, C., Brafman, R.I., Shimony, S.E.: Preference-based configuration of web page content. In: Proceedings of IJCAI 2001, pp. 1451–1456 (2001)
12. Gonzales, C., Perny, P., Queiroz, S.: Preference aggregation with graphical utility models. In: Proceedings of AAAI 2008, pp. 1037–1042 (2008)
13. Gottlob, G., Greco, G., Scarcello, F.: Pure Nash equilibria: Hard and easy games. J. Artif. Intell. Res. (JAIR) 24, 357–406 (2005)
14. Kearns, M., Littman, M.L., Singh, S.: Graphical models for game theory. In: Proceedings of UAI 2001 (2001)
15. Keeney, R.L., Raiffa, H.: Decision with Multiple Objectives: Preferences and Value Tradeoffs. Wiley, Chichester (1976)
16. Koller, D., Milch, B.: Multi-agent influence diagrams for representing and solving games. In: Proceedings of IJCAI 2001, pp. 1027–1034 (2001)
17. Lang, J., Xia, L.: Sequential composition of voting rules in multi-issue domains. In: Mathematical Social Sciences, pp. 304–324 (2009)
18. Prestwich, S., Rossi, F., Venable, B., Walsh, T.: Constrained CP nets. In: Italian Conference on Computational Logic (2004)

19. Rossi, F., Venable, K.B., Walsh, T.: mCP nets: Representing and reasoning with preferences of multiple agents. In: Proceedings of AAAI 2004, pp. 729–734 (2004)
20. Wilson, N.: Extending CP-nets with stronger conditional preference statements. In: Proceedings of AAAI 2004, pp. 735–741 (2004)
21. Wilson, N.: Efficient inference for expressive comparative preference languages. In: Proceedings of IJCAI 2009, pp. 961–966 (2009)
22. Xia, L., Conitzer, V., Lang, J.: Voting on multiattribute domains with cyclic preferential dependencies. In: Proceedings of AAAI 2008, pp. 202–207 (2008)