Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Models for Conditional Preferences as Extensions of CP-Nets

Cristina Cornelio
University of Padua
Padua, Italy
cornelio@math.unipd.it

Abstract

This paper presents two frameworks that generalize Conditional Preference networks (CP-nets). The first generalization is the LCP-theory, first order logic theory that provides a rich framework to express preferences. The the second generalization, the PCP-networks, is a probabilistic generalization of CP-nets that models conditional preferences with uncertainty.

1 Introduction

Preferences have an important role in AI and are involved in many situations of our life, therefore it is important to study how to model and reason with them. The ability of express preferences in an intuitive and compact way is essential, so many representations of preferences have been developed. We focus in particular on multi-attribute and structured frameworks. CP-nets [Boutilier *et al.*, 2004] are the main graphical model that represents conditional and qualitative preferences and there exists a lot of generalizations of this framework: the main are GCP-nets [Goldsmith *et al.*, 2008], CP-theories [Wilson, 2004] and comparative conditional theories [Wilson, 2009]. All these framework have a special syntax and semantic, need ad hoc algorithms and also they don't support probabilistic information.

With the purpose to unify all the conditional preference models, we provide a first generalization of CP-nets, that is, a first order logic framework: Logical Conditional Preference (LCP) theories. LCP-theories don't use special syntax or semantic, but they are based only on statements that are universally quantified first-order formulas. We provide an implementation of this new framework that have a computational complexity comparable with the state of the art procedures.

The Probabilistic Conditional Preference networks (PCP-nets) are the second generalization of CP-nets that we provide. This framework models conditional preferences with probabilistic uncertainty, maintaining the dependency structure employed by CP-nets. Given a PCP-net, we study deeply three reasoning tasks: aggregating a profile of agents (each one described by a CP-net) into a PCP-net, finding the outcome (a complete assignment to all the variables) that is most preferred by the agents or deciding whether one outcome is collectively preferred to another (a dominance query). Generally outcome optimization and dominance are performed directly and sequentially on the profile of CP-nets. We propose instead to aggregate the collection of CP-nets into a single structure, a PCP-net, on which we directly perform collective reasoning tasks: we can store or communicate to the agents a single, compact structure instead of a possibly large collection of CP-nets.

2 Logic formulation for Conditional Preference

LCP rules are universally quantified first-order formulas ($\forall c \rightarrow o \succ o'$) expressing preferences over outcomes o, o' for c a set of constraints (conjunction of equality atoms) over variables and constants.

Definition 1 *Given two distinct sets of n variables $\bar{X} = X_1, \cdots, X_n$ and $\bar{Y} = Y_1, \cdots, Y_n$ with domains $D_1, \cdots, D_n$ and c a* condition *(formula) with free variables in $\bar{X}$, $\bar{Y}$, consisting of a conjunction of equality formulas, a* **LCP-Rule** *is of the form: $\forall_{i\in\{1,\cdots,n\}} X_i, Y_i : D_i.\ c \rightarrow o \succ o'$, where o and o' are outcomes on $\bar{X}$ and $\bar{Y}$ respectively.*

Given a collection of LCP-rules $\mathcal{C}$, the corresponding **LCP-theory** *$\mathcal{L}[\mathcal{C}]$ has as axioms the formulas corresponding to the rules, and the transitivity axiom.*

LCP-Rules generate ground atomic formulas of the form $o_1 \cdots o_n \succ o'_1 \cdots o'_n$, each one is entailed for every choice of values for X_i (o_i) and Y_i (o'_i) that satisfies the condition described in c. It is important to notice that, in LCP rules, we directly represent an ordering on *outcomes* ($o \succ o'$) instead of on single features.

We provide algorithms for optimizing and comparing outcomes and for consistency checking. The computational complexity is higher since we are managing a more general model, but our results are comparable with the state of the art procedures.

We observe that CP-nets and the corresponding generalizations are special cases of LCP theories, but LCP theories also express other interesting kinds of preferences. For example, LCP theories can represent *recursive preferences*, that are statements that specify preferences conditioned on other preferences (e.g "If Alice prefers to go to the cinema then Bob prefers cinema too"). LCP theories can also represents minimal and maximal description of preferences (e.g. "I prefer everything rather than going to the cinema" or "I prefer to go to the cinema rather then any other option").

The powerful of the LCP formulation is to provide a formalization with no special syntax or semantics to represent preferences (contra the others frameworks [Boutilier *et al.*, 2004; Wilson, 2004; 2009; Bienvenu *et al.*, 2010]), generalization of the existing models, with a competitive computational complexity.

3 Probabilistic CP-nets (PCP-nets)

Definition 2 *A* **PCP-net (Probabilistic CP-net)** *[Cornelio* et al.*, 2013] is a directed graph where each node represents a variable (feature)* $\{X_1, \ldots, X_n\}$*, with finite domain. For each feature* X_i*, there is a set of* parent *features* $\mathrm{Pa}(X_i)$ *that can affect the preferences over the values of* X_i*. This defines a* dependency graph *in which each node* X_i *has edges from all features in* $\mathrm{Pa}(X_i)$*. Given this structural information, for each feature* X_i*, instead of giving, for each complete assignment on* $\mathrm{Pa}(X_i)$*, a total ordering over the domain of* X_i *(as in CP-nets), we give a probability distribution over the set of all the possible total orderings.*

PCP-nets are a strict generalization of CP-nets: when we restrict the probability distributions of a PCP-net in $\{0, 1\}$ we obtain the definition of CP-net. Thus, a PCP-net defines a probability distribution over a collection of CP-nets: the set of induced CP-nets that are CP-nets that can be obtained from the input PCP-net by choosing an ordering from each probability distribution over orderings.

Aggregation. We use PCP-nets in a multi-agent context to compactly represent a collection of CP-nets: we use probabilities to reconcile possibly conflicting preferences expressed by a group of agents. Given a profile of CP-nets (set of CP-nets on the same variables), there may not exists a PCP-net that induces exactly the same distribution over the initial profile of CP-nets. We define aggregation methods that work even in this case: we introduce and evaluate two aggregation methods [Cornelio *et al.*, 2015]: *PR* generates a PCP-net setting the probabilities in PCP-tables by adding the relative frequencies of the CP-nets in the profile with a particular configuration; *LS* minimizes the mean squared error between the probability distribution induced by the PCP-net over the input CP-nets and their relative frequency. Both methods are polynomial under certain conditions, but our theoretical and experimental results suggest that *PR* is more accurate.

Optimality. Finding an optimal outcome corresponds to find the outcome that best represents the preferences of the agents. We consider two notions of optimality for outcomes: *the most probable optimal outcome:* the outcome with the highest probability of being optimal; *the optimal outcome of the most probable induced CP-net:* the optimal outcome of the induced deterministic CP-net with the highest probability. Computing both of them takes polynomial time if the graph of the PCP-net has bounded width. Our experimental and theoretical results show that the optimal outcome of the most probable induced CP-net is the most accurate.

Dominance. Given a PCP-net, dominance returns the probability that one outcome is preferred to another. There are many results that show that this problem is hard [Bigot *et al.*, 2013; Boutilier *et al.*, 2004; Goldsmith *et al.*, 2008], thus we study algorithms to compute an approximation of this value. We define an approximation interval for the dominance value [Cornelio *et al.*, 2015], that can be computed in polynomial time. We observe that in the case of separable PCP-nets, the lower bound of this interval corresponds to the real value of dominance. This theoretical result suggests that the lower bound could be a good approximation of the real value of the dominance also for non-separable PCP-nets. The experimental evaluation confirms that the distance between the lower bound to the true dominance probability is generally very small.

4 Future work

There are many lines for future research: extend LCP-theories allowing hard constraints or include probabilities in LCP formulation allowing to manage uncertain scenarios with a richer framework for Conditional Preference. Another idea is to generalize CP-nets in a probabilistic framework using different probability theories: we generalized CP-theories using the Dempster-Shafer theory and we used the results obtained in [Moral and Wilson, 1994] to provide an efficient method that approximates dominance probability.

References

[Bienvenu *et al.*, 2010] M. Bienvenu, J. Lang, and N. Wilson. From Preference Logics to Preference Languages, and Back. In *Proc. of Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR 2010)*, 2010.

[Bigot *et al.*, 2013] D. Bigot, H. Fargier, J. Mengin, and B. Zanuttini. Probabilistic conditional preference networks. In *Proc. of the 29th International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.

[Boutilier *et al.*, 2004] C. Boutilier, R.I. Brafman, C. Domshlak, H.H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[Cornelio *et al.*, 2013] C. Cornelio, J. Goldsmith, N. Mattei, F. Rossi, and K.B. Venable. Updates and uncertainty in CP-nets. In *Proc. of the 26th Australasian Joint Conference on Artificial Intelligence (AUSAI)*, 2013.

[Cornelio *et al.*, 2015] C. Cornelio, J. Goldsmith, N. Mattei, F. Rossi, and K.B. Venable. Reasoning with PCP-nets in a multi-agent context. In *Proc. of the International Conference onAutonomous and Multiagent Systems (AAMAS)*, 2015.

[Goldsmith *et al.*, 2008] J. Goldsmith, J. Lang, M. Truszczynski, and N. Wilson. The computational complexity of dominance and consistency in CP-nets. *Journal of Artificial Intelligence Research*, 33(1):403–432, 2008.

[Moral and Wilson, 1994] S. Moral and N. Wilson. Markov chain monte-carlo algorithms for the calculation of dempster-shafer belief. In *Proc. of the Twelfth National Conference on Artificial Intelligence*, 1994.

[Wilson, 2004] N. Wilson. Extending CP-Nets with Stronger Conditional Preference Statements. In *Proc. of AAAI-04*, pages 735–741, 2004.

[Wilson, 2009] Nic Wilson. Efficient inference for expressive comparative preference languages. In *Proceedings of IJCAI-09*, 2009.