

NP-Completeness of Outcome Optimization for Partial CP-Nets

Keith Purrington and Edmund H. Durfee

Computer Science and Engineering
 University of Michigan
 Ann Arbor, Michigan 48109 USA
 <purringk, durfee>@umich.edu

CP-nets (Boutilier 2004) represent qualitative conditional preferences using a graphical model similar in concept to a belief network. A CP-net describes the preferences over the values of each problem variable, possibly conditioned on the values of other variables. For example, a user might describe her ideal dinner by expressing an (unconditional) preference for a meat entrée over a fish entrée, and a preference for wine that depends on the particular choice of entrée (red wine with meat, and white wine with fish). The CP-net for this diner would look like the one in Figure 1, with the exception that a preference would be expressed for the value of D . CP-nets make it easy to find the most preferred meal using a linear *forward sweep* algorithm. First, a meat entrée is selected based on the unconditional preference. Then, wine is selected to match the choice of entrée: red in this example.

CP-nets are an attractive representation of preferences in part because they facilitate preference elicitation without excessively burdening users. However, for scheduling problems where, for example, a user might prefer a morning meeting to an afternoon meeting, but have no more specific preferences, it can be unwieldy to require the user to express distinct preferences over each of the possible meeting times. Instead, it is much more natural if a user can express indifference between meetings at 9, 10, or 11 AM.

Partial CP-nets (Rossi 2004) address this concern by allowing the presence of *unranked variables* (variables about which an agent has expressed indifference) in a CP-net. This extension neatly addresses the representational issue but complicates the task of finding an optimal outcome. Rossi *et al.* show that the linear-time forward sweep algorithm for finding optimal outcomes in CP-nets translates directly to partial CP-nets when the agent is able to control the value of every variable. However, when external *evidence* fixes the values for some variables, the problem becomes much harder. The main contribution we report in this paper is a proof that for even a restrictive class of partial CP-nets, the problem of outcome optimization in the presence of evidence is NP-complete.

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Partial CP-Nets

The introduction of unranked variables has the main consequence of expanding the ways in which two outcomes that differ on the value of a single variable can be related. As in traditional CP-nets, when two outcomes differ on the value of one variable, then (since all else is equal) the outcome where that variable has its preferred value is preferred to the other outcome. However, if that variable is unranked, then preference over the outcomes is not so easily determined: even though neither value is directly preferred, one value might match better with variables whose preferred values are conditionally dependent upon it. For an unranked node I , a change to its value is worsening if it makes the values for all variables that depend on I less preferred. Similarly, a change is indifferent if it results in the value for each dependent variable being just as good. Finally, incomparable changes are defined to be those that are neither worsening nor indifferent.

The presence of evidence makes outcome optimization more difficult because such evidence effectively constrains the acceptable values for unranked nodes. For example, consider the very simple partial CP-net and induced outcome graph depicted in Figure 1. This network describes preferences for a food-wine pairing. For this individual, the particular choice of food is unimportant; her preference is rather that the food and wine go well together. This can be represented in several ways; here, the diner is indifferent as to choice of food, but expresses a preference for the type of wine that depends on what entrée is chosen.

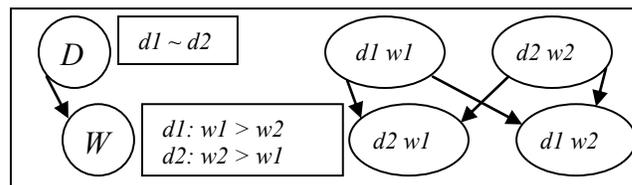


Figure 1. CP-net for diner indifferent to entrée choice.

Provided that the diner can choose all of the variable values, the forward sweep finds one of the most preferred meals. It can simply choose a value for D , and no matter

the choice, the CP-table for W provides a value that ensures a successful pairing. If, however, the restaurant is out of white wine (fixing the value of W to be w_2), a simple forward sweep is no longer guaranteed to find the optimal outcome. When the value of w_2 is fixed, the diner can still be maximally satisfied, but to do so with a forward sweep requires making the right arbitrary choice of red meat before discovering that the restaurant is out of white wine. To avoid the need for omniscience, a general optimization algorithm must allow for backtracking.

This small example can be solved with a *backward sweep* that uses the evidence to determine the proper value for the unranked variable. However, as the CP-net structure becomes more complex, a backward sweep encounters the same need for backtracking as needed by a forward sweep. In fact, we show that even for a very restricted class of problems, making assignments to the unranked nodes so as to ensure that the value of each evidence node is preferred is NP-complete.

NP-Completeness of Outcome Optimization

We summarize here our proof of the NP-completeness of outcome optimization for a very specialized class of partial CP-nets. We consider partial CP-nets with a bipartite graph structure $\langle I, E, D \rangle$, where I is a set of unranked nodes, E is a set of evidence nodes, and D is a set of directed edges from nodes in I to nodes in E . These restricted CP-nets look like the one in Figure 2. Finally, we restrict the variables in these specialized networks to be binary-valued.

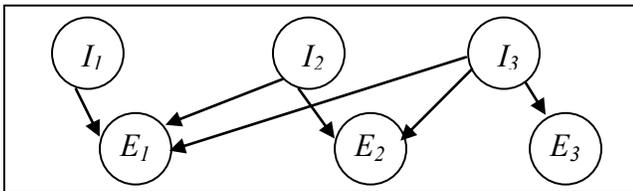


Figure 2. Network showing bipartite structure

We show hardness by reducing 3SAT to the problem of finding the optimal outcome in such a graph, by mapping 3SAT clauses to evidence nodes. By carefully choosing the parent nodes and the CP-tables for the evidence nodes, we ensure that each evidence variable is satisfied IFF the corresponding 3SAT clause is satisfied.

Consider a particular 3SAT clause, say $(I_1 \text{ OR } I_2 \text{ OR } \sim I_3)$. To translate this to our CP-net, we create an evidence node E , with three parents I_1 , I_2 , and I_3 . To construct the CP-table for E , say that a value of e (rather than $\sim e$) satisfies the node. Then, because E has three binary-valued parents, there are eight entries in its CP-table. We fill in the table so that e is preferred to $\sim e$ in exactly the rows that correspond to satisfying assignments to the SAT clause. For each evidence node, seven of the eight possible parent assignments will satisfy that evidence. In this case, only for $(\sim i_1, \sim i_2, i_3)$ do we set $\sim e$ preferable to e . Performing the

same procedure to create an evidence node for each 3SAT clause completes the reduction.

It is straightforward to show that the problem is also contained in NP. A solution certificate can be produced by nondeterministically choosing an assignment for the unranked nodes. Then, this certificate can be checked to see if it simultaneously satisfies the evidence nodes in polynomial time via simple inspection of the CP-table for each evidence node.

Algorithmic Approximation

While evidence-constrained partial CP-net optimization is hard in general, a backward sweep algorithm as discussed in conjunction with the diner example remains optimal in certain special cases, namely when each unranked node has at most one path to an evidence node. In such networks, a three-step algorithm finds optimal outcomes in polynomial time. The three algorithmic steps consist of a forward sweep to assign variables that are independent of the unranked nodes, followed by a backward sweep that works back from the evidence nodes to determine optimal settings for the unranked nodes, and finally, an additional forward sweep to assign the variables that are dependent on the unranked nodes. In addition to its provable optimality on problems with the special structure, when the algorithm is applied to general problems, it greatly outperforms a GSAT-style random-walk algorithm, finding slightly better solutions, while operating many times faster. Details of the algorithm and its empirical evaluation, as well as a more fully detailed NP-completeness proof, are available from the authors upon request.

Acknowledgments

The work reported in this paper was supported, in part, by the National Science Foundation under grant IIS-0534280. We would like to thank Jim Boerkoel and the anonymous reviewers for valuable feedback on this work.

References

- Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H.H., and Poole, D. 2003. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. In *Journal of Artificial Intelligence Research (JAIR)* 21 (2004), 135-191.
- Rossi, F., Venable, K.B., and Walsh, T. mCP Nets: Representing and Reasoning with Preferences of Multiple Agents. In: *Procs. of the 19th Conf. on Artificial Intelligence*, LNCS 749, AAAI Press (2004) 729-734