Optimal Decision Making with CP-nets and PCP-nets

(Extended Abstract)

Sujoy Sikdar
Rensselaer Polytechnic Institute
Computer Science Department
sikdas@rpi.edu

Sibel Adalı
Rensselaer Polytechnic Institute
Computer Science Department
sibel@cs.rpi.edu

Lirong Xia
Rensselaer Polytechnic Institute
Computer Science Department
xial@cs.rpi.edu

ABSTRACT

Probabilistic conditional preference networks (PCP-nets) are a generalization of CP-nets for compactly representing preferences over multi-attribute domains. We introduce the notion of a loss function whose inputs are a CP-net and an outcome. We focus on the optimal decision-making problem for acyclic and cyclic CP-nets and PCP-nets. Our motivations are three-fold: (1) our framework naturally extends to allow reasoning on cyclic CP-nets and PCP-nets for full generality, (2) in the multi-agent setting, we place no restriction on agents' preferences structure and voting rules under our framework have desirable axiomatic properties, (3) we generalize several previous approaches to finding the optimum outcome in individual and multi-agent contexts. We characterize the computational complexity of computing the loss of a given outcome and computing the outcomes with minimum loss for three natural loss functions: 0-1 *loss, neighborhood loss, and global loss*. While the optimal decision is NP-hard to compute for many cases, we give a polynomial-time algorithm for computing the optimal decision for tree-structured PCP-nets and profiles of CP-net preferences with a shared dependency structure, w.r.t. neighborhood loss function.

1. INTRODUCTION

Many decision-making problems involve choosing an optimal outcome from a *multi-attribute domain* where the alternatives are characterized by $p \geq 1$ variables and each variable corresponds to an attribute of the outcome. In combinatorial voting there are p issues, and the alternatives correspond to the decisions made on each issue. The goal is to make an optimal (joint) decision for an agent or a group of agents with preferences over the alternatives. However, since the number of outcomes in a multi-attribute domain is exponentially large, it is impractical for the agents to express preferences as a full ranking over all outcomes.

A popular practical solution is to use a compact preference language to represent agents' preferences. Perhaps the most commonly used language for agents to represent their preferences over multi-attribute domains are CP-nets (conditional preference networks) [2]. In a CP-net, an agent can specify her local preferences over any attribute given the values of some other attributes (called its parents). Agents' preferences are expressed in terms of *ceteris paribus statements* of the form: "I prefer red wine to white wine, ceteris paribus, given that meat is served as the main dish."

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

For a single agent whose preferences are represented by a CP-net, a natural optimization objective is to identify *undominated* outcomes [3]. Informally, an outcome is undominated if no other outcome is preferred over it. The problem of computing undominated outcomes is well studied in the CP-net literature. For acyclic CP-nets (CP-nets with acyclic dependency graphs), an undominated outcome always exists and is unique [2]. However, when we allow cyclic dependencies, undominated outcomes can be hard to compute [3, 9].

Recently, probabilistic conditional preference networks (PCP-nets) have been introduced as a natural generalization of CP-nets [1, 7]. A PCP-net can be used to represent a single agent's uncertain preferences over a set of CP-nets, or a preference profile of multiple CP-nets [8]. Given an acyclic PCP-net, [7] provides a polynomial-time algorithm for computing the outcome that is undominated with the highest probability. Despite this promising first step in decision making with PCP-nets, the optimal decision making problem for PCP-nets remains largely open. In particular, is there any other sensible and more quantitative optimality criterion beyond "being undominated" that we may consider for CP-nets as well as PCP-nets? If so, how can we compute them? We take a decision-theoretic approach by modeling the optimality of an outcome by a *loss function*, whose inputs are an outcome (an assignment of values to attributes) and a single (acyclic or cyclic) CP-net.

However, much of the existing work [11, 20, 18, 14, 19, 15, 5, 12] focuses on certain special cases with rather severe restrictions on agents' preferences such as allowing only profiles with acyclic CP-nets, and dependencies that are compatible with a common order on the issues (O-legality). We design a new class of voting rules under our loss minimization framework that are characterized by a loss function which takes as input any profile of CP-net preferences and outputs a set of loss minimizing outcomes.

2. PRELIMINARIES

Let $I = \{X_1, ..., X_p\}$ be a finite set of p variables with finite domains $D(X_i)$. Let $\mathcal{L}(D(X_i))$ denote the set of all linear orders over $D(X_i)$. For ease of presentation, *we will assume that all variables are binary in this paper*. An assignment (or outcome) $\vec{d}$ is a vector in $\Pi_{i \leq p} D(X_i)$. We use d_i to denote the value of X_i in $\vec{d}$, and d_{-i} to denote the values of all other variables.

DEFINITION 1. *[2] A CP-net C over the set of variables I is given by two components (i) a directed graph $G = (I, E)$ called the* dependency *graph, and (ii) for each variable X_i, there is a* conditional preference table *$CPT(X_i)$ that contains a linear order $\succ^i_{C,\vec{u}}$ over $D(X_i)$ for each valuation $\vec{u}$ of the parents of X_i (denoted $Pa(X_i)$) in G.*

When G is (a)cyclic we say that C is a (a)cyclic CP-net.

The partial order $\succ_C$ induced by a CP-net C over the set of all possible assignments $\Pi_{i \le p} D(X_i)$ is the transitive closure of $\{(a_i, \vec{u}, \vec{z}) \succ (b_i, \vec{u}, \vec{z})\} : i \le p; a_i, b_i \in D(X_i); \vec{u} \in D(Pa(X_i)); \vec{z} \in D(-(Pa(X_i) \cup \{X_i\}))\}$. A CP-net is said to be consistent if $\succ_C$ is asymmetric. Acyclic CP-nets are consistent but cyclic CP-nets are not necessarily consistent.

DEFINITION 2 (WEAK AND STRICT DOMINANCE). *An assignment $\vec{a}$* weakly dominates *$\vec{b}$ if $\vec{a} \succ_C \vec{b}$. An assignment $\vec{a}$* strictly dominates *$\vec{b}$ if $\vec{a} \succ_C \vec{b}$ and $\vec{b} \not\succ_C \vec{a}$.*

DEFINITION 3. *A PCP-net [1, 7] Q over the set of variables I is given by (i) a directed graph $G = (I, E)$, and (ii) for each variable X_i, there is a* probabilistic conditional preference table *$PCPT(X_i)$ that contains a probability distribution $f^i_{Q,\vec{u}}$ over $\mathcal{L}(D(X_i))$ for each valuation $\vec{u}$ of the parents of X_i in G.*

Loss Functions. In this paper we will focus on three loss functions. Each loss function L takes a CP-net C and an assignment $\vec{d}$ as inputs and outputs a real number $L(C, \vec{d})$.

DEFINITION 4. *The 0-1 loss function is defined as*

$$L_{0-1}(C, \vec{d}) = \begin{cases} 1 & \text{if there exists } \vec{d'} \text{ such that } \vec{d'} \succ_C \vec{d}, \\ 0 & \text{otherwise} \end{cases}$$

That is, the 0-1 loss function takes the value 0 if and only if $\vec{d}$ is not weakly dominated by any other assignment in C.

DEFINITION 5. *The neighborhood loss function is defined as $L_N(C, \vec{d}) = |\{\vec{d'} : \exists i : d'_i \succ_C d_i \text{ and } d'_{-i} = d_{-i}\}|$.*

That is, the neighborhood loss of $\vec{d}$ in C is the number of $\vec{d}$'s neighbors that can be obtained by a single improving flip from $\vec{d}$ in C.

DEFINITION 6. *The global loss function is defined as $L_G(C, \vec{d}) = |\{\vec{d'} : \vec{d'} \succ_C \vec{d}, \text{and } \vec{d} \not\succ_C \vec{d'}|$.*

That is, the global loss of $\vec{d}$ in C is the total number of assignments that strictly dominate $\vec{d}$ in C.

We now formally define the decision problem of computing the loss of an assignment w.r.t. a loss function.

DEFINITION 7 (L-LOSS). *Given a PCP-net Q, a loss function L, a decision $\vec{d}$, and a number $k \in \mathbb{R}$, in L-LOSS we are asked to compute whether $L(Q, \vec{d}) \le k$.*

Optimal Decision Problem. We define the decision problem of computing optimal assignments L-OPTDECISION as follows.

DEFINITION 8 (L-OPTDECISION). *Given a PCP-net Q, a loss function L, and a number $k \in \mathbb{R}$, does there exist an assignment $\vec{d}$ such that $L(Q, \vec{d}) \le k$?*

Loss Minimizing Voting Rules. We develop a new class of voting rules characterized by a loss function. Given a loss function L, we define the voting rule r_L to be the function that maps a CP-net profile to a decision that minimizes the total loss i.e. $r_L(P) = \arg\min_{\vec{d}} L(P, \vec{d})$, where $L(P, \vec{d}) = \sum_{i=1}^{n} L(P_i, \vec{d})$. We define the decision problem of computing optimal joint decisions under this setting as L-OPTJOINTDECISION:

DEFINITION 9 (L-OPTJOINTDECISION). *Given a profile P, a collection of CP-net preferences, a loss function L, and a number $k \in \mathbb{R}$, does there exist an assignment $\vec{d}$ such that $L(P, \vec{d}) \le k$?*

3. DISCUSSION OF MAIN RESULTS

One might be tempted to believe that PCP-nets are so complicated that all problems are hard to compute. This is not true. As we can see in Table 1, computing LOSS w.r.t. L_{0-1} and L_N can be done in polynomial time for PCP-nets. Another false belief could be that for the same loss function, LOSS is easier than OPTDECISION (or vice versa). Neither is true by comparing Table 2(a) and Table 1. L_G-LOSS is coNP-hard but L_G-OPTDECISION is in P for acyclic CP-nets. L_N-LOSS is in P but L_N-OPTDECISION is NP-complete for cyclic CP-nets. While it is hard to compute the optimal outcomes w.r.t. all three loss functions (Table 2), for tree-structured PCP-nets, we have a polynomial-time algorithm to compute the optimal outcome. Similarly, while it is, hard to compute the optimal outcomes w.r.t. L_{0-1} for acyclic PCP-nets, a simple polynomial time algorithm allows us to compute the optimal outcome for a profile of acyclic CP-nets. Finally, every voting rule under our framework satisfies *anonymity, category-wise neutrality, consistency and weak monotonicity*.

Table 1: Complexity of L-LOSS w.r.t. acyclic and cyclic CP-nets. The complexity remains unchanged for the case of acyclic and cyclic PCP-nets.

<table>
<tr><th>Loss fn.</th><th>Acyclic</th><th>Cyclic</th></tr>
<tr><td>L_{0-1}</td><td rowspan="2">P (trivial)</td><td>P</td></tr>
<tr><td>L_N</td><td>P</td></tr>
<tr><td>L_G</td><td>coNPH</td><td>coNPH</td></tr>
</table>

Table 2: Complexity of L-OPTDECISION w.r.t. acyclic and cyclic CP-nets and PCP-nets.

<table>
<tr><th rowspan="2">Loss fn.</th><th colspan="2">CP-nets</th><th colspan="2">PCP-nets</th></tr>
<tr><th>Acyclic</th><th>Cyclic</th><th>Acyclic</th><th>Cyclic</th></tr>
<tr><td>$L_{0,1}$</td><td rowspan="3">P [2]</td><td rowspan="2">NPC</td><td>NPC, P for trees [7]</td><td>NPC [7]</td></tr>
<tr><td>L_N</td><td>NPH, P for trees</td><td>NPH</td></tr>
<tr><td>L_G</td><td>P</td><td colspan="2">coNPH</td></tr>
</table>

Table 3: Complexity of L-OPTJOINTDECISION w.r.t. profiles of acyclic and cyclic CP-nets.

<table>
<tr><th>Loss fn.</th><th>Acyclic</th><th>Cyclic</th></tr>
<tr><td>L_{0-1}</td><td>P</td><td>NPC</td></tr>
<tr><td>L_N</td><td colspan="2">NPC, P for shared tree-structured dependency graph.</td></tr>
<tr><td>L_G</td><td colspan="2">coNPH</td></tr>
</table>

Acknowledgments

We thank the anonymous reviewers for their many inputs and suggestions. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053 (the ARL Network Science CTA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Lirong Xia acknowledges National Science Foundation under grant IIS-1453542 for support.

REFERENCES

[1] D. Bigot, B. Zanuttini, H. Fargier, and J. Mengin. Probabilistic Conditional Preference Networks. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, 2013.

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

[3] R. I. Brafman and Y. Dimopoulos. Extended Semantics and Optimization Algorithms for CP-Networks. *Computational Intelligence*, 20(2):218—245, 2004.

[4] F. Brandt, V. Conitzer, and U. Endriss. Computational social choice. In G. Weiss, editor, *Multiagent Systems*. MIT Press, 2013.

[5] V. Conitzer, J. Lang, and L. Xia. Hypercubewise preference aggregation in multi-issue domains. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, Barcelona, Catalonia, Spain, 2011.

[6] V. Conitzer and L. Xia. Paradoxes of Multiple Elections: An Approximation Approach. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 179–187, Rome, Italy, 2012.

[7] C. Cornelio, J. Goldsmith, N. Mattei, F. Rossi, and K. B. Venable. Updates and uncertainty in CP-nets. In *Advances in Artificial Intelligence*, pages 301–312. Springer, 2013.

[8] C. Cornelio, U. Grandi, J. Goldsmith, N. Mattei, F. Rossi, and K. B. Venable. Voting with CP-nets using a probabilistic preference structure. In *Proceedings of COMSOC*, 2014.

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

[10] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.

[11] J. Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1366–1371, Hyderabad, India, 2007.

[12] J. Lang, J. Mengin, and L. Xia. Aggregating conditionally lexicographic preferences on multi-issue domains. In *Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming (CP)*, pages 973–987, 2012.

[13] J. Lang and L. Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[14] M. Li, Q. B. Vo, and R. Kowalczyk. An Efficient Procedure for Collective Decision-making with CP-nets. In *Proceedings of the 19th European Conference on Artificial Intelligence*, pages 375–380, 2010.

[15] M. Li, Q. B. Vo, and R. Kowalczyk. Majority-rule-based preference aggregation on multi-attribute domains with structured preferences. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2011.

[16] E. Mackin and L. Xia. Allocating indivisible items in categorized domains. In *Proceedings of IJCAI-2016*, 2016.

[17] F. Rossi, K. B. Venable, and T. Walsh. mCP nets: Representing and reasoning with preferences of multiple agents. In *Proceedings of AAAI-04*, pages 729–734, 2004.

[18] L. Xia, V. Conitzer, and J. Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 202–207, Chicago, IL, USA, 2008.

[19] L. Xia, V. Conitzer, and J. Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 399–406, Toronto, Canada, 2010.

[20] L. Xia, J. Lang, and M. Ying. Strongly decomposable voting rules on multiattribute domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 776–781, Vancouver, BC, Canada, 2007.