An Efficient Majority-Rule-Based Approach for Collective Decision Making with CP-Nets

Minyi Li
Swinburne Univ. of Technology
myli@swin.edu.au

Quoc Bao Vo
Swinburne Univ. of Technology
BVo@swin.edu.au

Ryszard Kowalczyk
Swinburne Univ. of Technology
RKowalczyk@swin.edu.au

Abstract

This paper addresses the problem of collective decision making in the case where the agents' preferences are represented by CP-nets (Conditional Preference Networks). Most existing works either do not consider the computational issues, or depend on a strong assumption that all the agents share a common preferential independence structure. To this end, this paper proposes an efficient approach, called `CDMCP` (Collective Decision Making with CP-nets), for aggregating multiple agents' preferences according to majority rule. The proposed approach allows the agents to have different preferential independence structures and is computationally efficient.

Introduction

In many real world scenarios, we need to represent and reason about the simultaneous preferences of several agents, and to aggregate such preferences (Rossi, Venable, and Walsh 2004). Classical decision theory considers typically agents as having preferences over the outcome space, and the preferences are usually mathematically represented by utility functions. However, in many situations, the utility-based preference elicitation is not user-friendly (Boutilier et al. 2004). In this paper, we investigate the theory of CP-nets as a formal model for representing and reasoning with the agents' preferences. It specifies individual preference relations in a relatively compact, intuitive, and structured manner (Boutilier et al. 2004). Most existing works on CP-nets focus on individual preference reasoning, including outcome optimization and comparison (Boutilier et al. 1999; 2004). However, many real world scenarios require aggregating multiple agents' preferences. Rossi *et al.* (Rossi, Venable, and Walsh 2004) propose various voting semantics for aggregating multiple agents' CP-nets. However, they do not address computational issues. Given the fact that answering dominance query[1] for an arbitrary CP-net is PSPACE-complete (Goldsmith et al. 2008). Outcome optimization with multiple agents furthermore requires searching the entire outcome space, and making dominance queries to each agent on each pair of outcomes. This problem is likely not a member of NP or coNP, but even harder. Consequently, their approach may not be applicable when the number of possible outcomes from which to choose is large, especially when the set of outcomes has a *combinatorial structure*[2] (Chevaleyre et al. 2008). Lang (Xia, Lang, and Ying 2007; Lang and Xia 2009) reconsiders voting and aggregation rules in the case where the agents' preferences have a common preferential independence structure. The author addresses the decompositions with a voting rule following a linear order over variables. Xia *et al.* (Xia, Conitzer, and Lang 2008) extended the setting in (Xia, Lang, and Ying 2007) such that the CP-nets do not necessarily follow a common order. However, their proposed approach does not come with any practical algorithm for computing the set of majority-optimal outcomes.

To this end, this paper addresses the above drawbacks, proposing an efficient approach, called `CDMCP` (Collective Decision Making with CP-nets), for aggregating multiple agents' preferences according to majority rule. The proposed approach allows the agents to have different preferential independence structures and is computationally efficient.

Background

CP-nets overview

Let $\mathbf{V} = \{X_1, \ldots, X_m\}$ be a set of m variables. For each $X_k \in \mathbf{V}$, $D(X_k)$ is the *value domain* of X_k. A variable X_k is *binary* if $D(X_k) = \{x_k, \bar{x}_k\}$. If $\{x_k, \bar{x}_k\}$ is a binary domain of X_k, then $x_k = \neg\bar{x}_k$; $\bar{x}_k = \neg x_k$. If $\mathbf{X} = \left\{X_{i_1}, \ldots, X_{i_p}\right\} \subseteq \mathbf{V}$, with $i_1 < \cdots < i_p$ then $D(\mathbf{X})$ denotes $D(X_{i_1}) \times \cdots \times D\left(X_{i_p}\right)$. The assignments of variable values to $\mathbf{X}$ are denoted by $\mathbf{x}$, $\mathbf{x}'$ etc., and represented by concatenating the values of the variables. For instance, if $\mathbf{X} = \{X_1, X_2, X_3\}$, an assignment $\mathbf{x} = x_1\bar{x}_2x_3$ assigns x_1 to X_1, $\bar{x}_2$ to X_2 and x_3 to X_3. If $\mathbf{X} = \mathbf{V}$, $\mathbf{x}$ is a *complete assignment*; otherwise $\mathbf{x}$ is called a *partial assignment*. For an assignment $\mathbf{x}$, we denote by $\mathbf{x}[X_k]$ the value $x_k \in D(X_k)$ assigned to variable X_k by that assignment. We also allow logic operations between the value assignments to binary variables. For instance, $x_1\bar{x}_2 = x_1 \wedge \bar{x}_2 = (X_1 = x_1) \wedge (X_2 = \bar{x}_2)$. That is, x_1 is $True$

[1]A dominance query, given a pair of outcomes, asks whether one outcome dominates another.

[2]The number of all possible outcome is exponential in the number of variables.

and x_2 is $False$ (i.e. $\neg\bar{x}_2 = x_2$). If $\mathbf{p} = x_1\bar{x}_2$ and $\mathbf{q} = x_3$, then $\mathbf{p} \vee \mathbf{q} = (x_1\bar{x}_2) \vee x_3 = ((X_1 = x_1) \wedge (X_2 = \bar{x}_2)) \vee (X_3 = x_3)$.

Let $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ be nonempty sets that partition $\mathbf{V}$ and $\succ$ a preference relation over $D(\mathbf{V})$. $\mathbf{X}$ is *(conditionally) preferentially independent* of $\mathbf{Y}$ given $\mathbf{Z}$ if and only if, for all $\mathbf{x}, \mathbf{x}' \in D(\mathbf{X})$, $\mathbf{y}, \mathbf{y}' \in D(\mathbf{Y})$, $\mathbf{z} \in D(\mathbf{Z})$:

$$\mathbf{xyz} \succ \mathbf{x}'\mathbf{yz} \text{ iff } \mathbf{xy}'\mathbf{z} \succ \mathbf{x}'\mathbf{y}'\mathbf{z}$$

A CP-net $\mathcal{N}$ (Boutilier et al. 2004) over $\mathbf{V}$ is an annotated directed graph $\mathcal{G}$ over $X_1, \ldots, X_m$, in which nodes stand for the problem variables. Each node X_k is annotated with a conditional preference table $CPT(X_k)$, which associates a total order $\succ^{X_k|Pa(X_k)=\mathbf{p}}$ with each instantiation $\mathbf{p}$ of X_k's parents $Pa(X_k)$. For instance, let $\mathbf{V} = \{X_1, X_2, X_3\}$, all three being binary, and assume that preference of a given agent over all possible outcomes can be defined by a CP-net whose structural part is the directed acyclic graph $\mathcal{G} = \{(X_1, X_2), (X_1, X_3), (X_2, X_3)\}$; this means that the agent's preference over the values of X_1 is unconditional, preference over the values of X_2 (resp. X_3) is fully determined given the values of X_1 (resp. the values of X_1 and X_2). The conditional preference statements contained in these tables are written with the usual notation, that is, $x_1\bar{x}_2$ means that when x_1 is true and x_2 is false then $X_3 = x_3$ is preferred to $X_3 = \bar{x}_3$.

Majority rule

Given a set of outcomes, we need to aggregate multiple agents' preference and decide on one as the final outcome. Rossi *et al.* (Rossi, Venable, and Walsh 2004) propose a *Majority Voting Semantics* for aggregating multiple agents' preferences which are represented by CP-nets. Given two outcomes o_1 and o_2, let $\mathbf{S}_\succ$, $\mathbf{S}_\prec$, $\mathbf{S}_\approx$, $\mathbf{S}_\bowtie$ be the sets of agents who say, respectively, that $o_1 \succ o_2$, $o_1 \prec o_2$, $o_1 \approx o_2$ (indifferent), and $o_1 \bowtie o_2$ (incomparable).

Definition 1 (Majority). *Given two outcomes o_1 and o_2, o_1 is majority better than o_2 (written as $o_1 \succ_{maj} o_2$) if and only if there is a majority of agents who prefer o_1 to o_2 (i.e. $|\mathbf{S}_\succ| > |\mathbf{S}_\prec| + |\mathbf{S}_\bowtie|$). Two outcomes are majority incomparable if and only if they are not ordered in either way. Moreover, an outcome is majority-optimal if and only if no other outcome is majority better than that outcome.*

When the preference ordering of each agent over the outcome space is linear (total order), the majority-optimal outcome is unique if it exists. However, due to the incompleteness of preference relations induced by CP-nets, there may be more than one majority-optimal outcomes in the case where the agents' preferences are represented by CP-nets. Note that any pair of outcomes in the set of majority-optimal outcomes are majority incomparable.

Assume that there are n agents making decisions over a set of m binary variables. To test whether an outcome is majority-optimal we need to compare the given outcome with all other outcomes ($O(2^m)$) in all CP-nets (n). Consequently, testing majority-optimality has the complexity of $O(n2^m)$ and there are not lower bounds (Rossi, Venable, and Walsh 2004)[3]. Moreover, outcome optimization with multiple agents' CP-nets is even more challenge. We need to compare all outcomes (2^m) to all other outcomes (2^m) in all CP-nets(n). That means, finding the set of majority-optimal outcomes has the following complexity: $O(n2^{2m})$. In the following section, we present an approach for finding the set of possible majority-optimal outcomes more efficiently by reducing the search space into a small subset of outcomes.

The proposed approach

In this section, we present our proposed approach `CDMCP` for collective decision making with multiple CP-nets based on majority rule. The propose approach generates a set of variable value constraints. We then employ a SAT solver to compute the set of possible majority-optimal outcomes which satisfy all these constraints. By doing so, the brute-force optimality checking can be done only on the remaining outcomes, and thus the proposed approach is computationally efficient.

Assume that there are a set of n agents $\mathbf{A} = \{A_1, \ldots, A_n\}$ making decisions over a set of m variables $\mathbf{V} = \{X_1, \ldots, X_m\}$. The preference of each agent A_i is captured by a CP-net $\mathcal{N}_i$.

Step 1: For each variable X_i with a binary value domain $\{x_i, \bar{x}_i\}$, each agent A_j's has a conditional preference table $CPT_j(X_i)$ stating the conditional preference on the values of variable X_i with each instantiation of X_i's parents $Pa_j(X_i)$. For each agent A_j, we separate these instantiations of $Pa_j(X_i)$ into the following two categories.

- $\mathbf{P}_{A_j}^{x_i \succ \bar{x}_i} = \left\{\mathbf{p} \in D\left(Pa_j(X_i)\right) \mid x_i \succ_{A_j}^{X_i|\mathbf{p}} \bar{x}_i\right\}$.
  Let $\rho_{A_j}^{x_i \succ \bar{x}_i} \triangleq \bigvee_{\mathbf{p} \in \mathbf{P}_{A_j}^{x_i \succ \bar{x}_i}} \mathbf{p}$, thus $x_i \succ_{A_j}^{X_i|\rho_{A_j}^{x_i \succ \bar{x}_i}} \bar{x}_i$.
- $\mathbf{P}_{A_j}^{\bar{x}_i \succ x_i} = \left\{\mathbf{p} \in D\left(Pa_j\left(X_j\right)\right) \mid \bar{x}_i \succ_{A_j}^{X_i|\mathbf{p}} x_i\right\}$.
  Let $\rho_{A_j}^{\bar{x}_i \succ x_i} \triangleq \bigvee_{\mathbf{p} \in \mathbf{P}_{A_j}^{\bar{x}_i \succ x_i}} \mathbf{p}$, thus $\bar{x}_i \succ_{A_j}^{X_i|\rho_{A_j}^{\bar{x}_i \succ x_i}} x_i$.

Let $\mathbf{u}$ be an assignment that satisfies $\rho_{A_j}^{x_i \succ \bar{x}_i}$ (resp. $\rho_{A_j}^{\bar{x}_i \succ x_i}$), such that $\mathbf{U}$ is the set of variables assigned by $\mathbf{u}$ and $\mathbf{W} = \mathbf{V} - \mathbf{U} - \{X_i\}$, then $\mathbf{u}x_i\mathbf{w} \succ_{A_j} \mathbf{u}\bar{x}_i\mathbf{w}$ (resp. $\mathbf{u}\bar{x}_i\mathbf{w} \succ_{A_j} \mathbf{u}x_i\mathbf{w}$) for all $\mathbf{w} \in D(\mathbf{W})$. If a given agent has unconditional preference over a variable X_i: $x_i \succ_{A_j}^{X_i} \bar{x}_i$ (resp. $\bar{x}_i \succ_{A_j}^{X_i} x_i$), that means the condition $\rho_{A_j}^{x_i \succ \bar{x}_i}$ (resp. $\rho_{A_j}^{\bar{x}_i \succ x_i}$) is always $True$. On the other hand, $\rho_{A_j}^{\bar{x}_i \succ x_i}$ (resp. $\rho_{A_j}^{x_i \succ \bar{x}_i}$) is always $False$.

Let $k = (n+1)/2$. Given two outcomes o_1 and o_2, we say that o_1 is majority better than o_2 (written as $o_1 \succ_{maj} o_2$) if and only if $|\mathbf{S}_\succ| \geq k$. Moreover, there will be a set of $\binom{n}{k}$ combinations of agents that satisfy the majority requirement, denoted by *qComs*; and each combination *com* (*com* $\in$ *qComs*) is a distinct subset of $\mathbf{A}$ which contains k agents. A combination of agents *com* $\in$ *qComs* prefers x_i to $\bar{x}_i$ ceteris paribus under a certain instantiation $\mathbf{u}$ such

[3]Note that in this paper, all the agents are considering the same set of variables

that $\bigwedge_{A_j \in com} \rho_{A_j}^{x_i > \bar{x}_i}$ is $True$. Let $\varphi_{com}^{x_i > \bar{x}_i}$ denote $\bigwedge_{A_j \in com} \rho_{A_j}^{x_i > \bar{x}_i}$, then:

$$x_i \succ_{com}^{X_i|\varphi_{com}^{x_i > \bar{x}_i}} \bar{x}_i \tag{1}$$

Similarly, let $\varphi_{com}^{\bar{x}_i > x_i}$ denote $\bigwedge_{A_j \in com} \rho_{A_j}^{\bar{x}_i > x_i}$, then:

$$\bar{x}_i \succ_{com}^{X_i|\varphi_{com}^{\bar{x}_i > x_i}} x_i \tag{2}$$

Let $\mathbf{u}$ be an assignment that satisfies $\varphi_{com}^{x_i > \bar{x}_i}$ (resp. $\varphi_{com}^{\bar{x}_i > x_i}$), such that $\mathbf{U}$ is the set of variables assigned by $\mathbf{u}$ and $\mathbf{W} = \mathbf{V} - \mathbf{U} - \{X_i\}$, then $\mathbf{u}x_i\mathbf{w} \succ_{com} \mathbf{u}\bar{x}_i\mathbf{w}$ (resp. $\mathbf{u}\bar{x}_i\mathbf{w} \succ_{com} \mathbf{u}x_i\mathbf{w}$) for all $\mathbf{w} \in D(\mathbf{W})$.

Based on majority rule, if there exists a combination $com \in qComs$, such that $\varphi_{com}^{x_i > \bar{x}_i}$ is $True$, then a majority number of agents prefer x_i to $\bar{x}_i$ ceteris paribus. Let $\gamma_{maj}^{x_i > \bar{x}_i}$ denote $\bigvee_{com \in qComs} \varphi_{com}^{x_i > \bar{x}_i}$, then:

$$x_i \succ_{maj}^{X_i|\gamma_{maj}^{x_i > \bar{x}_i}} \bar{x}_i \tag{3}$$

Similarly, let $\gamma_{maj}^{\bar{x}_i > x_i}$ denote $\bigvee_{com \in qComs} \varphi_{com}^{\bar{x}_i > x_i}$, then:

$$\bar{x}_i \succ_{maj}^{X_i|\gamma_{maj}^{\bar{x}_i > x_i}} x_i \tag{4}$$

Let $\mathbf{u}$ be an assignment that satisfies $\gamma_{maj}^{\bar{x}_i > x_i}$ (resp. $\gamma_{maj}^{x_i > \bar{x}_i}$), such that $\mathbf{U}$ is the set of variables assigned by $\mathbf{u}$ and $\mathbf{W} = \mathbf{V} - \mathbf{U} - \{X_i\}$, then $\mathbf{u}\bar{x}_i\mathbf{w} \succ_{maj} \mathbf{u}x_i\mathbf{w}$ (resp. $\mathbf{u}x_i\mathbf{w} \succ_{maj} \mathbf{u}\bar{x}_i\mathbf{w}$) for all $\mathbf{w} \in D(\mathbf{W})$.

Consequently, for each $X_i \in \mathbf{V}$, a majority-optimal outcome o^* must satisfy the following constraint S_i:

$$S_i = \left(\gamma^{x_i > \bar{x}_i} \Rightarrow x_i\right) \wedge \left(\gamma^{\bar{x}_i > x_i} \Rightarrow \bar{x}\right)$$

Let $S = \bigwedge_{X_i \in \mathbf{V}} S_i$. Then a possible majority-optimal outcome must be the model of S. We then employ a SAT solver to this constraint satisfaction problem. The SAT solver returns a set of models of S, which are the possible majority-optimal outcomes. Note that any other outcome that is not a model of S would not be majority-optimal. Consequently, we can only test the majority-optimality of the outcomes in this set.

Conclusion and future work

In this paper, we have introduced an efficient approach to compute the set of majority-optimal outcomes from a collection of CP-nets. There are not many works for aggregating multiple agents' preferences represented by CP-nets. Unlike previous work where voters' preferences are required to satisfy some restrictive conditions on the dependence graph (such as the existence of a common acyclic graph to all agent), the proposed method applies to all profiles.

We are now planning to explore more powerful variants such as TCP-nets for representing agents' preferences and investigate techniques to aggregate such preferences.

References

Boutilier, C.; Brafman, R. I.; Hoos, H. H.; and Poole, D. 1999. Reasoning with conditional ceteris paribus preference statements. 71–80. Morgan Kaufmann Publishers.

Boutilier, C.; Brafman, R. I.; Hoos, H. H.; and Poole, D. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research* 21:2004.

Chevaleyre, Y.; Endriss, U.; Lang, J.; and Maudet, N. 2008. Preference handling in combinatorial domains: From ai to social choice. *AI Magazine, Special Issue on Preferences* 29(4):37–46.

Goldsmith, J.; Lang, J.; Truszczynski, M.; and Wilson, N. 2008. The computational complexity of dominance and consistency in CP-nets. *J. Artif. Int. Res.* 33(1):403–432.

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

Rossi, F.; Venable, K. B.; and Walsh, T. 2004. mCP nets: Representing and reasoning with preferences of multiple agents. In McGuinness, D. L., and Ferguson, G., eds., *AAAI*, 729–734. AAAI Press / The MIT Press.

Xia, L.; Conitzer, V.; and Lang, J. 2008. Voting on multiattribute domains with cyclic preferential dependencies. In *AAAI'08: Proceedings of the 23rd national conference on Artificial intelligence*, 202–207. AAAI Press.

Xia, L.; Lang, J.; and Ying, M. 2007. Strongly decomposable voting rules on multiattribute domains. In *AAAI'07: Proceedings of the 22nd national conference on Artificial intelligence*, 776–781. AAAI Press.