Efficient Penalty Scoring Functions for Group Decision-making with TCP-nets

(Extended Abstract)

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

Quoc Bao Vo
Swinburne University of Technology
BVO@swin.edu.au

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

ABSTRACT

This paper studies the problem of collective decision-making in combinatorial domain where the agents' preferences are represented by qualitative models with TCP-nets (Tradeoffs-enhanced Conditional Preference Network). The features of TCP-nets enable us to easily encode human preferences and the relative importance between the decision variables; however, many group decision-making methods require numerical measures of degrees of desirability of alternative outcomes. To permit a natural way for preference elicitation while providing quantitative comparisons between outcomes, we present a computationally efficient approach that compiles individual TCP-nets into ordinal penalty scoring functions. After the individual penalty scores are computed, we further define a collective penalty scoring function to aggregate multiple agents' preferences.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Algorithms, Design

Keywords

Group decision-making; TCP-nets; Penalty scoring function

1. INTRODUCTION

In many real world scenarios, we need to represent and reason about the simultaneous preferences of multiple agents [4]. In this paper, we investigate the theory of TCP-nets (Tradeoffs-enhanced Conditional Preference Network) [2], a variant of CP-net (Conditional Preference Network) [1], as a formal model for representing and reasoning about the agents' preferences. We present an approach that compiles an individual TCP-net into an ordinal penalty scoring function. The proposed approach preserves all strict preference ordering induced by the original TCP-net and provides a numerical measure of desirability of alternative outcomes. Moreover, it provides an easy way for preferential comparisons. After the individual penalty scores of each agent is built, then the individual penalty scores are aggregated into a normalized collective penalty scoring function modelling the preferences of a group of agents.

Cite as: Efficient Penalty Scoring Functions for Group Decision-making with TCP-nets (Extended Abstract), Minyi Li, Quoc Bao Vo and Ryszard Kowalczyk, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1073-1074.

2. TCP-NETS

A TCP-net (Tradeoffs-enhanced Conditional Preference Network) $\mathcal{N}$ [2] is a preference-representation structure that extends the CP-net [1] by incorporating the *relative importance* between variables. The nodes of a TCP-net are the domain variables. There are three sets of arcs between variables: cp, i and ci. cp denotes a set of *directed* cp-arcs (cp stands for conditional preference). A cp-arc $\langle\overrightarrow{X,Y}\rangle$ is in $\mathcal{N}$ iff the preferences over the values of Y depend on the actual value of X; we called X is a parent variable of Y. Each variable Y is then annotated with a conditional preference table $CPT(Y)$, which associates a total order $\succ^{Y|\mathbf{u}}$ with each instantiation $\mathbf{u}$ of Y's parents $Pa(Y)$, i.e. $\mathbf{u} \in D(Pa(Y))$. i is a set of directed i-arcs (where i stands for *importance*). An i-arc $\langle\overrightarrow{X,Y}\rangle$ is in $\mathcal{N}$ iff X is unconditionally more important than Y, i.e., $X \triangleright Y$. ci is a set of *undirected* ci-arcs (where ci stands for conditional importance). A ci-arc (X, Y) is in $\mathcal{N}$ iff we have $\mathcal{RI}(X, Y|\mathbf{Z})$ for some $\mathbf{Z} \subseteq \mathbf{V} - \{X, Y\}$ and $\mathbf{Z}$ is called the *selector set* of (X, Y). We denote the *selector set* of a ci-arc $\gamma = (X, Y)$ by $\mathcal{S}(\gamma)$ and the union of the selector set in a TCP-net $\mathcal{N}$ by $\mathcal{S}(\mathcal{N})$. Each ci-arc $\gamma = (X, Y)$ is associated with a conditional importance table $CIT(\gamma)$ from every instantiation of $\mathbf{s} \in D(\mathcal{S}(\gamma))$ to the orders over the set $\{X, Y\}$. A TCP-net in which the sets i and ci (and therefore, the conditional importance tables) are empty, is also a CP-net. In this paper, we make the classical assumption that each agent j's TCP-nets $\mathcal{N}_j$ is conditionally acyclic[1].

3. INDIVIDUAL PREFERENCE

Our work of individual preference approximation is based on the work of Domshlak *et al.* [3], which provides a numerical approximation for acyclic CP-nets using weighted soft constraints. In this paper, we go one step further by incorporating the relative importance information among pairs of variables and introduce an ordinal penalty scoring function as a numerical approximation not only for acyclic CP-nets, but also for conditionally acyclic TCP-nets. In broad terms, given a conditionally acyclic TCP-net, we generate a penalty scoring function representing that TCP-net in the following steps. First, we assign an *importance weight* to each variable based on the structure of the given TCP-net. Next, a penalty scoring function is defined based on penalty analysis. As to examine the structure induced by a TCP-net, we recall the following notion of the dependency graph of a TCP-net [2]:

DEFINITION 1 (DEPENDENCY GRAPH). *The dependency graph $\mathcal{N}^*$ of a TCP-net $\mathcal{N}$ contains all the nodes and arcs of $\mathcal{N}$, and for every* ci*-arc $\gamma = (X, Y)$ in $\mathcal{N}$ and every variable $Z \in \mathcal{S}(\gamma)$, $\mathcal{N}^*$*

[1] We refer to [2] for the formal definition of conditionally acyclic TCP-nets.

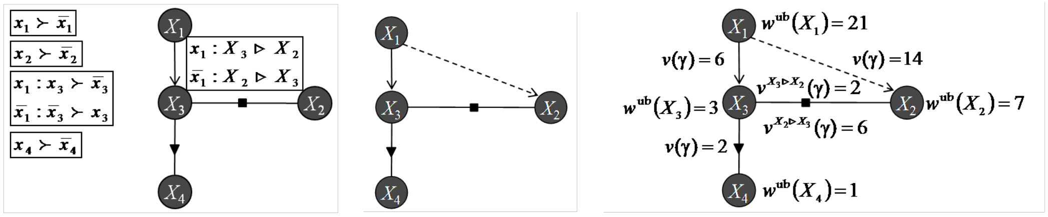

(a) TCP-net (b) Dependency graph (c) Arc value

Figure 1: An example of TCP-net, its dependency graph and arc values

contains a directed sci-*arc* $\langle\overrightarrow{Z,X}\rangle$ *(resp.* $\langle\overrightarrow{Z,Y}\rangle$*), if there is no arc between* Z *and* X *(resp.* Z *and* Y*) in* $\mathcal{N}$*. We denote* sci *as the set of* sci-*arcs in* $\mathcal{N}^*$.

Figure 1(b) shows the dependency graph of the CP-net in Figure 1(a). For a variable X in $\mathcal{N}^*$, we call the variables Y s.t. $\langle\overrightarrow{X,Y}\rangle \in \mathsf{cp} \cup \mathsf{sci}$ as the dependants of X. For a variable X, let $|D(X)|$ be the domain size of X and thus there are $|D(X)|$ degrees of penalties. Without loss of generality, we assume the degree of penalties of a variable X range between 0 and $|D(X)|-1$, that is, $d_1 = 0, \ldots, d_{|D(X)|} = |D(X)| - 1$. For the TCP-nets in Figure 1(a), since all variables are binary, there are only two degrees of penalties, i.e., $d_1 = 0$ and $d_2 = 1$. For a variable X, consider a preference ordering over the value of X given an instantiation of X's parents, let the rank of the most preferred value of X be 0 and the rank of the least preferred value of X be $|D(X)| - 1$. Consequently, given an outcome o, the degree of penalty of a variable X in o is the rank of the value $o[X]$ in the preference ordering over X given the parent context $\mathbf{u} = o[Pa(X)]$. We denote by d_X^o ($d_X^o \in \left\{d_1, \ldots, d_{|D(X)|}\right\}$) the degree of penalty of X with respect to o. For instance, consider a variable X such that $D(X) = \{x, x', x''\}$. Assume that, under a parent context $\mathbf{u} = o[Pa(X)]$ assigned by an outcome o, $x \succ x' \succ x''$. Hence, if $o[X] = x$, then $d_X^o = d_1 = 0$; if $o[X] = x'$, then $d_X^o = d_2 = 1$; if $o[X] = x''$, then $d_X^o = d_3 = 2$.

We then analyse the *importance weights* of variables in a TCP-net. We first assign the value to each arc in the dependency graph of the given TCP-net, then, we analyse the importance weight of a variable X in a particular outcome o, denoted by $w^o(X)$, by considering *(i)* the values of the directed cp-, i- and sci-arcs $\langle\overrightarrow{X,Y}\rangle$ that originate at X; and *(ii)* the values of the ci-arcs $\gamma = (X,Y) \in \mathsf{ci}$ s.t. $X \rhd Y$ given $\mathbf{z} = o[\mathcal{S}(\gamma)]$. We denote the value of an arc γ where $\gamma \in \mathsf{cp} \cup \mathsf{sci} \cup \mathsf{i}$ by $v(\gamma)$; and the value of an arc $\gamma = (X,Y) \in \mathsf{ci}$ under the condition that $X \rhd Y$ (resp. $Y \rhd Z$) by $v^{X \rhd Y}(\gamma)$ (resp. $v^{Y \rhd X}(\gamma)$). Moreover, as the importance weight of a variable X is *context-dependent*, when we assign the value to an arc γ, we consider the upper bound weight of X that γ points to. The upper bound weight of a variable X, denoted by $w^{\mathsf{ub}}(X_1)$, is computed under the assumption that for all ci-arc $(X,Y) \in \mathsf{ci}$, X is contextually more important than Y. Figure 1(c) shows an example of assignments to the arc values and upper bound weights of variables for the given dependency graph in Figure 1(b).

Given a TCP-net $\mathcal{N}$ and an outcome o, the penalty of a variable X in o is the degree of penalty of X in o, i.e. d_X^o, multiplied by the importance weight of X in o, i.e. $w^o(X)$. Then we can analyse the penalty score of an outcome by considering the sum of the penalties of variables in that outcome: $\forall o \in O$, $\mathtt{pen}(o) = \sum_{X \in \mathbf{V}} w^o(X) \cdot d_X^o$

4. COLLECTIVE PREFERENCE

After the individual penalty scores are computed independently, these penalty scores are aggregated into a normalized collective penalty scoring function that best conveys the preferences of the group of the agents.

Definition 2 (Collective penalty scoring function). *Given a set of conditionally acyclic TCP-nets* $\mathbf{N} = \{\mathcal{N}_1, \ldots, \mathcal{N}_n\}$*, the collective penalty scoring function* P *mapping from* O *to* $[0, +\infty]$ *is defined by:*

$$\forall o \in O,\ \mathtt{P}(o) = \diamond\{pen_i(o) \mid i = 1, \ldots, n\} \tag{1}$$

where $\diamond$ *is a function that satisfies non-decreasingness for each of its argument and commutativity.*

As discussed in [4], the most natural choices for $\diamond$ are `sum` and `max`. `sum` is a *utilitarian* aggregation operator, stating that the collective penalty score of an outcome is the sum of the penalty scores of the agents in the group. On the other hand, `max` states that the maximum penalty score among all the agents should be considered. Thus, the `max` aggregation operator corresponds to the *egalitarian social welfare*.

5. FUTURE WORK

. In this paper, we have studied the problem of group decision-making with TCP-nets (Tradeoffs-enhanced Conditional Preference Network). Based on the previous work, we have gone one step further by incorporating the relative importance relation among pairs of variables and introduced an individual penalty scoring function as a numerical approximation not only for acyclic CP-nets, but also for conditionally acyclic TCP-nets.

Nonetheless, the present work is only applicable for conditionally acyclic TCP-nets. The investigation of techniques to deal with cyclic preferences need to be further explored.

6. ACKNOWLEDGEMENTS

This work is partially supported by the ARC Discovery Grant DP0987380.

7. REFERENCES

[1] C. Boutilier, R. I. Brafman, 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*, 21:2004, 2003.

[2] R. I. Brafman, C. Domshlak, and S. E. Shimony. On graphical modeling of preference and importance. *Journal of AI Research*, 25:2006, 2006.

[3] C. Domshlak, S. D. Prestwich, F. Rossi, K. B. Venable, and T. Walsh. Hard and soft constraints for reasoning about qualitative conditional preferences. *J. Heuristics*, 12(4-5):263–285, 2006.

[4] C. Lafage and J. Lang. Logical representation of preferences for group decision making. In *KR*, pages 457–468, 2000.