Influence and aggregation of preferences over combinatorial domains

(Extended Abstract)

Nicolas Maudet
LIP6, UPMC
Paris, France
nicolas.maudet@lip6.fr

Maria Silvia Pini
Univ. of Padova, Italy
mpini@math.unipd.it

Francesca Rossi
Univ. of Padova, Italy
frossi@math.unipd.it

Kristen Brent Venable
Univ. of Padova, Italy
kvenable@math.unipd.it

ABSTRACT

In a multi-agent context where a set of agents declares their preferences over a common set of candidates, it is often the case that agents may influence each others. Recent work has modelled the influence phenomenon in the case of voting over a single issue. Here we generalize this model to account for preferences over combinatorially structured domains including several issues. When agents express their preferences as CP-nets, we show how to model influence functions and how to aggregate preferences by interleaving voting and influence convergence.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Theory, Algorithms

Keywords

Preference aggregation, influence, combinatorial domains

1. INTRODUCTION

Often a set of agents needs to select a common decision from a set of possible decisions, over which they express their preferences, and such a decision set has a combinatorial structure. That is, it can be seen as the combination of certain issues, where each issue has a set of possible instances. Consider for example a car: usually it is not seen as a single item, but as a combination of features, such as its engine, its shape, its color, and its cost. Each of these features has some possible instances, and a car is the combination of such feature instances. If a family needs to buy a new car, each family member may have his own opinion about each feature of a car, and the task is to choose the

car that best fits the preferences of everybody. But suppose the mother knows well the CO2 emissions of the different cars: her preference regarding the engine may affect the one of his son who is concerned by the carbon footprint. In other words, agents may *influence* each other, leading their preferences to be modified accordingly.

The concept of influence has been widely studied in psychology, economics, sociology, and mathematics. Recent work has modelled the influence phenomenon in the case of taking a decision over a single binary issue [2]. Under this iterative model of influence, we may pass from state to state until stability holds, or we may also not converge.

Here we generalize this model to account for preferences over combinatorially structured domains including several issues. Complex influence statements may be represented, *e.g.* influences which depend on the context ("if my daughter prefers the yellow color for the car, I will follow her; otherwise I will stick to my inclinations") or which may involve different features of different agents ("If my wife and my son prefer the small car, then I would prefer the green color").

Usually preferences over combinatorially structured domains are expressed compactly, otherwise too much space would be needed to rank all possible alternatives. CP-nets are a successful framework that allows one to do this [1]. They exploit the independence among some features to give conditional preferences over small subsets of them. CP-nets have already been considered in a multi-agent setting [5, 4]. Here we incorporate influences among agents.

2. MODELLING PREFERENCES AND INFLUENCES

We assume each agent expresses its preferences over the candidates via an acyclic CP-net [1]. CP-nets are sets of *conditional preference statements* (cp-statements) each stating a total order over the values of a variable (say X), possibly depending on each combination of values of a set of other variables (say $X_1, \ldots, X_n$). X is said the dependent variable and $X_1, \ldots, X_n$ are the parents of X. Acyclic CP-nets are CP-nets where the dependency graph (with arcs from parents to dependent variables) does not have cycles.

We also assume that the dependency graphs of such CP-nets must all be compatible with a linear order O over the features: for each voter, the preference over a feature is in-

dependent of features following it in O (O-legality in [3]). A *profile* models the *initial inclination* of all agents (their opinions over the candidates before they are influenced by each other) as a collection of n such acyclic CP-nets over the m features.

To model influences, we use conditional influence statements. A *conditional influence statement* (ci-statement) on variable X has the form

$$O(X_1), \ldots, O(X_k) :: o(X)$$

where $o(Y)$ is an ordering over the values of variable Y, for $Y \in \{X_1, \ldots, X_n, X\}$. Variables $X_1, \ldots X_k$ are the influencing variables and variable X is the influenced variable.

A ci-table is a collection of ci-statements with the same influencing and influenced variables, and containing at most one ci-statement for each ordering of the influencing variables. An *I-profile* is a triple (P, O, S), where P is a profile, O is an ordering over the m features of the profile, and S is a set of ci-tables. We assume that the ci-tables of an I-profile must be such that each variable can be influenced only by variables in her level or in earlier levels, but not in the same ci-statement. Thus, ci-arcs in an I-profile can create cycles only within variables of the same level.

EXAMPLE 1. *There are three agents (thus three CP-nets, all compatible with the ordering $X \succ Y$), and two binary features: X and Y, with values, respectively, x and $\overline{x}$, and y and $\overline{y}$. The I-profile has six variables denoted by X_1, X_2, X_3, Y_1, Y_2, and Y_3. Each variable X_i (resp., Y_i), with $i \in \{1, 2, 3\}$, has two values denoted by x_i and $\overline{x_i}$ (resp., y_i and $\overline{y_i}$). Note that cp-statements are denoted by single-line arrows while ci-statements are denoted by doubled-line arrows. As it can be seen, agent 3 is influenced on feature X by agent 2.*

There is a very useful relationship between ci-statements and cp-statements:

THEOREM 1. *Given an influence function f, consider the set of cp-statements N corresponding the ci-statements $ci(f)$. Then the undominated outcomes of N coincide with the stable states of f.*

While this result allows for a very simple integration of ci- and cp-statements in the same profile, it is important to still distinguish between the initial inclinations (cp-statements) and the influences (ci-statements). In fact, influences modify the initial inclination by overriding the preferences, but the opposite does not hold.

3. AGGREGATING PREFERENCES

We propose a way to aggregate the preferences contained in an I-profile, while taking into account the influence functions. The method we propose includes three main phases:

- *Influence iteration within one level*: For each feature, we consider the influences among different variables modelling this feature. An iterative algorithm is used: it takes all variables regarding the same feature and starts with the assignment corresponding to the initial inclination. The output is a single state. Either the algorithm, by iteratively applying the influences, ended up in a stable state; or it detected a cycle and used a subroutine to select nevertheless a single state.

- *Propagation from one level to the next one*: Once the variables of a certain level have been fixed to some values, we propagate to the next level by considering the ci- and cp-statements that go from this level to the next one. As influence overrides preference, we first look at the ci-tables and set the inclination of the influenced variables according to such tables. For the variables whose inclination has not been determined after this step, their inclination will be determined by their cp-tables. After this, we are ready to handle the next level as we did for the first one, since all of its variables are now subject only to influence functions.

- *Preference aggregation*: Since at each level we obtain a possibly different value for the variables modelling the same feature, we may either aggregate at each level (LA, *Level Aggregation*) or only at the end of the procedure (FA, *Final Aggregation*) when each agent has its most preferred candidate. Under LA (using majority since variables are binary) we assign the same value to all variables. Then we propagate such a choice to the next level and start again with an influence iteration. Under FA, we leave the variable values in each level as they are after the influence iteration and proceed until all levels have been handled. At this point, we have a most preferred candidate for each agent, and we can obtain a winning candidate by any voting rule that needs the top choices, such as plurality.

The two approaches may yield different results (this can observed on our example, where the winner is $\langle X = x, Y = y \rangle$ under LA and $(X = \overline{x}, Y = \overline{y})$ under FA. However, the choice of the ordering O does not matter as far as the winner is concerned, no matter if we use LA or FA.

4. REFERENCES

[1] 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. *JAIR*, 21:135–191, 2004.

[2] M. Grabisch and A. Rusinowska. Iterating influence between players in a social network. In *Proceedings of the 16th Coalition Theory Network Workshop*, 2011.

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

[4] K. Purrington and E. H. Durfee. Making social choices from individuals' cp-nets. In *AAMAS*, page 179, 2007.

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