Applied Mathematics and Computation 332 (2018) 96–104

Contents lists available at ScienceDirect

Applied Mathematics and Computation

journal homepage: www.elsevier.com/locate/amc

Strategic advice for decision-making under conflict based on observed behaviour

Amanda Garcia[a,*], Amer Obeidi[b], Keith W. Hipel[a]

[a] Department of Systems Design Engineering, University of Waterloo, Waterloo, Canada
[b] Department of Management Science and Department of Systems Design Engineering, University of Waterloo, Waterloo, Canada

ARTICLE INFO

Keywords:
Decision-making
Graph Model for Conflict Resolution
Inverse engineering
Game theory
Preferences

ABSTRACT

An improvement in the inverse engineering of preferences approach for the Graph Model for Conflict Resolution is introduced. In addition to providing decision-makers and analysts with up-to-date preference information about opponents, the methodology is now equipped with an Advice function which enriches the decision-making process by providing important information regarding potential moves. Decision-makers who use the methods introduced in this paper are provided with the expected value of each of their possible moves, with the probability of the opponent's next response, and with the opponent reachable states. This insightful information helps establish an accurate picture of the conflict situation and in so doing, aids stakeholders in making strategic decisions.

1. Introduction

The study of human behaviour in conflict seeks to model the dynamics of choices and interactions between the decision-makers (DMs) in the dispute. This field of study has evolved from von Newmann and Morgenstern's static game theoretic models [1] to include contributions from complexity science [2], network analysis [3,4], and statistical physics [5,6]. Within some of these models, preference elicitation is a key step.

Both "forward" and "inverse" approaches have been developed to assist analysts determine a DM's preferences over a list of objects or outcomes. Generally speaking, the former attempt to construct preference rankings based on attributes, values, or judgments provided by the DM; common methods in this vein include the analytic hierarchy process (AHP) [7,8], ELimination and Choice Expressing the REality (ELECTRE) [9,10] and value-focused thinking [11]. Inverse approaches, on the other hand, work "backwards" from observed behaviours to a reward function which contains preference information. These approaches commonly use Markov Decision Processes (MDPs) or Partially Observable MDPs (POMDPs) to infer a DM's reward function given an observed behaviour [12–14]. Such techniques are designed to solve problems in inverse reinforcement learning: one can only observe behaviours from an expert and must use these observations to learn to perform these same tasks.

Preference elicitation is a critical step in the Graph Model for Conflict Resolution (Graph Model) methodology [15,16]; the model relies on preference input from all of the DMs in order to produce its analysis results. Within the Graph Model paradigm, "forward" preference elicitation is akin to eliciting ordinal preferences from DMs. Several techniques can be used to perform this task, including option weighting, option prioritizing, and direct ranking. More recently, work has been done

* Corresponding author.
E-mail address: a9garcia@uwaterloo.ca (A. Garcia).

https://doi.org/10.1016/j.amc.2018.03.031

to integrate AHP within the Graph Model [17] and to extend option prioritization to non-standard preference structures. In this case, the analyst is working closely with all of the DMs involved in the conflict and, in doing so, is able to elicit the necessary preference information to carry out a stability analysis.

Preferences themselves have been studied from many angles; a common practice is to use different preferences structures to allow for uncertainty (e.g. fuzzy preferences [18]) or to capture an element of DM behaviour (e.g., attitudes [19]). Other Graph Model approaches can partially or fully bypass preferences altogether; the inverse approach to the Graph Model for Conflict Resolution starts from a desired outcome and provides one or more DMs with the preferences necessary to realise it [20], while Sakakibara et. al. provide a method to perform a Graph Model analysis with incomplete information [21].

Adopting an inverse approach for preference elicitation within the Graph Model is a relatively new endeavour [22]. This procedure has several advantages: first, there is a lesser need for direct consultation between DMs and analysts since a DM's preferences can be inferred from their behaviour alone. This is particularly relevant for adversarial situations in which DMs may not wish to communicate their preferences to analysts in order to preserve a strategic advantage. This characteristic is also helpful in the study of historical conflicts for which documentation might be sparse. Rather than relying on an analyst's best guess or on an approximation of DM preferences, empirical data are used to generate insights.

Second, the dynamic nature of the algorithms allows for ongoing analysis which adapts to the latest changes in the conflict. This contrasts with the static approaches of the Graph Model which typically selects a single point in time for analysis.

Finally, the inverse approach can be used to provide relevant advice to DMs about what to do next. The Advice function which is the focus of this paper is designed to provide key information such as expected value and probability of occurrence. Advice is updated as the conflict progresses, and DMs can use the information provided by opponent moves to determine their best counter-moves and strategies in real time. The goal is not to replace a DM in the decision-making process, but to supply enriching information for reaching more informed decisions.

This paper is organized as follows: Section 2 sets the groundwork for the Graph Model and for inverse engineering preferences; Section 3 details the workings of the advice function, provides pseudocode, and discusses algorithmic complexity; Section 4 applies the algorithms to a small case study; and Section 5 provides future research avenues and conclusions.

2. The Graph Model and inverse engineering preferences

This section lays the theoretical groundwork both for the Graph Model and the inverse engineering methodologies on which the advice function is based. The overview provided here is non-technical; those details can be found in [15,23] and [22].

2.1. The Graph Model

The Graph Model is a methodology based in game theory used to analyse strategic interactions among several DMs [15,16,23]. The methodology requires DMs, the set of options under each of their control, and each DM's preferences over the set of conflict states. Model analysis reveals the stability or instability of conflict states as well as equilibrium states. A variety of stability concepts is used to model different types of human behaviour under conflict. The various stability concepts take into account levels of foresight and risk aversion, amongst others, to provide rich analytical results.

Once the relevant DMs have been identified, an option-based approach [24] can be used to determine the set of states in the conflict. An option tableau outlines the options for each DM. Options are assigned a value of 1 if they have been selected by the DM and 0 if they have not been chosen. It is common to display DMs and their options in the option form pioneered by Howard [24].

A DM's choice of one or more options under its control constitutes a strategy for that DM; conflict states are determined by combining strategies from each of the DMs. Given a state, each DM can unilaterally move away from it by changing one or more of the options under its control; such a move is called a unilateral move (UM).

A conflict with m options across all DMs will have a total of 2^m mathematically possible states; however, not all these states are feasible. States may be logically or preferentially infeasible for one or more DMs; such states are usually dropped from the conflict model in order to streamline analysis. Logically infeasible (e.g. mutually exclusive) states are always removed from the analysis since they do not represent a possible state of affairs. Preferentially infeasible states (e.g. it is highly unlikely that a DM would choose a particular option) must be removed with care as they can produce false equilibria [25].

Next, each DM's preferences over the set of states are recorded. Preference information for DMs is commonly elicited via interviews with DMs or through rigorous research when DMs are not available. Once all of the preference information has been collected, the analysis stage can begin. With a DM's preference information in hand, one can, given a state, determine a DM's set of unilateral improvements (UIs) from that state. Unilateral improvements are simply unilateral moves which result in a move preferred state for the mover.

In the analysis stage, states are classified as unstable (i.e., not likely to last) or stable (i.e., likely to last) for one or more DMs under a particular stability concept. Stability concepts mathematically express variations of human behaviour under conflict. Using a variety of stability concepts allows one to see how different DMs might react to a given situation. Depending