Reliability Engineering and System Safety 000 (2017) 1–9

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

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

# A graphical method for simplifying Bayesian games

Peter A. Thwaites[a,*], Jim Q. Smith[b]

[a] School of Mathematics, University of Leeds, LS2 9JT, Leeds, United Kingdom
[b] Department of Statistics, University of Warwick, Coventry, CV4 7AL, and Alan Turing Institute, United Kingdom

ARTICLE INFO

*Keywords:*
Adversarial risk
Bayesian game theory
Chain event graph
Decision tree
Influence diagram
Parsimony

ABSTRACT

If the influence diagram (ID) depicting a Bayesian game is common knowledge to its players then additional assumptions may allow the players to make use of its embodied irrelevance statements. They can then use these to discover a simpler game which still embodies both their optimal decision policies. However the impact of this result has been rather limited because many common Bayesian games do not exhibit sufficient symmetry to be fully and efficiently represented by an ID. The tree-based chain event graph (CEG) has been developed specifically for such asymmetric problems. By using these graphs rational players can make analogous deductions, assuming the topology of the CEG as common knowledge. In this paper we describe these powerful new techniques and illustrate them through an example modelling a game played between a government department and the provider of a website designed to radicalise vulnerable people.



## 1. Introduction

There are two principal conceptual difficulties in applying results from Bayesian game theory in a number of domains. Firstly, whilst it might be plausible for a player to know the broad structure of an opponent's utility function when that opponent is subjective expected utility maximizing (SEUM), for a player to also believe that she knows the *exact* quantitative form of that utility function or the precise formulation of the distribution of its attributes is less plausible. Secondly, as for example Nau [1] has pointed out, however compelling our beliefs are that an opponent's rationality *should* induce her to be SEUM, in practice most people simply are not. So any application of a theory which starts with this assumption is hazardous. These issues induced Kadane and Larkey [2] to suggest giving up on the rationality hypothesis entirely and instead modelling the opponent simply in terms of her past behaviour.

However others have perservered with rationality modelling by addressing these real modelling challenges more qualitatively. For example Smith [3] suggested a way to address the first difficulty described above. We can continue to model successfully provided that the conditional independences associated with various hypotheses and the attributes of each player's utility function are common knowledge, but we do not need that the players know the quantitative forms of others' inputs. This framework developed from methods for simplifying influence diagrams (IDs) [4,5], described first in [6] and then [7]. When players are all SEUM, substantive conclusions can sometimes be made concerning those aspects of the problem upon which a rational opponent's decision rules might depend. This in turn allows players to determine ever simpler forms for their own optimal decision rules. So models can be built which at least respect some of the structural implications of rationality hypotheses **before** being embellished with further structure gleaned from behavioural data, or the bold assumption that an opponent's quantitative preferences and beliefs can be fully quantified by everyone. Even the second criticism of a Bayesian approach outlined above is at least partially addressed, since the methods need only certain **structural** implications of SEUM to be valid, not that all players are SEUM.

In this way game theory can therefore be used not to fully specify the quantitative form of a competitive domain but simply to provide hypotheses about the likely *dependence* structure that rationality assumptions might imply for such models. These models can then be embellished with further historical quantitative information using the conventional Bayesian paradigm.

It has been possible to demonstrate the efficacy of the approach when modelling certain rather domain-specific applications, but it has proved rather limited in scope [3,8]. One problem is that the structure of many games cannot be fully and effectively represented by an ID (see for example [9–11]). Usually the underlying game tree is highly asymmetric and so the symmetries necessary for an encompassing and parsimonious ID representation of the game are not present. This is one characteristic of the types of games that we consider in this paper; we discuss other important attributes in the following paragraphs.

* Corresponding author.
*E-mail addresses:* P.A.Thwaites@leeds.ac.uk (P.A. Thwaites), J.Q.Smith@warwick.ac.uk (J.Q. Smith).

http://dx.doi.org/10.1016/j.ress.2017.05.012
Available online xxx
0951-8320/

Harsanyi [12] considered games where the players are uncertain about some or all of the following – the other players' utility functions, the strategies available to the other players, and the information other players have about the game. In the games considered in this paper each player holds a body of common knowledge – the exact form of other players' utility functions is unknown, but the variables these functions depend upon (a feature of the conditional independence structure of the game) are known; the strategies available to the other players are known; and what information is known to the other players is also known. So essentially the games we consider are ones where the "structure" is common knowledge, but the exact values of other players' utilities and the probability distributions of some chance variables are not.

In contrast to Harsanyi, Banks et al. in [13] state that *Game theory needs the defenders to know the attackers' utilities and probabilities, and the attackers to know the defenders' utilities and probabilities, and for both to know that these are common knowledge.* We do not agree that these are absolute requirements, but it is certainly true that a player cannot *solve* a game to her satisfaction unless she has some values for her opponents' utilities and probabilities. So in our games, players assign subjective probabilities to their *unknowns* and estimate values of their opponents' utility functions. Each player's utilities depend not only on the strategies chosen by the various players, but also on chance.

We take a decision-theoretic approach to Bayesian Game theory. Our games are sequential (typically with players acting alternately, and with chance variables interspersed between the players' actions). The standard description for such a game is *Extensive Form Bayesian Game with Chance moves*; they are generally expressed as a *game tree* (or as an ID [3] or MAID – multi-agent influence diagram [14]).

Asymmetric games, as described above, are being played with increasing frequency wherever large constitutional organisations (governments, police forces etc.) are at risk from or attempting to combat criminal or anticonstitutional organisations or networks. An example is described below, taken from this area, probably less familiar than games in a commercial context.

Governments and police play a game with groups trying to influence or radicalise susceptible individuals. These *radicalisers* often attempt to influence vulnerable people via the web. The government strategy here can be thought of as a combination of *prevention* and *pursuit*: if a website is easily accessible then it might be best just to shut it down; if it is difficult to access, then perhaps it is better to monitor, collect information and then act to scare vulnerable people sufficiently so that they do not get involved with any anticonstitutional group. But when should the government act? There is a trade-off here between frustrating a number of attempts to radicalise vulnerable people, and bringing down a whole anticonstitutional group (with the attached risks of failure and of exposing more susceptible individuals to malign influence for a longer period of time). The decisions available to the radicalisers are similar; the asymmetry of the game arises from the fact that different decisions by both players lead to very different collections of possible futures.

The Chain Event Graph (CEG) was introduced in 2008 [15] for the modelling of probabilistic problems whose underlying trees exhibit a high degree of asymmetry. It provides a platform from which to deduce dependence relationships between variables directly from the graph's topology. CEGs have principally been used for learning/model selection (see for example [16,17]), but also in two areas of interest to us in this paper – causal analysis (see for example [18,19]), and also decision analysis [20] where the semantics of the CEG can be extended to provide algorithms which allow users to discover minimal sets of variables needed to fully specify an SEUM decision rule. In 2015 it was realised that CEGs include Acyclic Probabilistic Finite Automata (APFAs) as a special case [21].

In this paper we demonstrate how it is often possible to use causal CEGs to deduce (from appropriate qualitative assumptions) a simpler representation of a two person game. To retain plausibility we assume only the qualitative structure of the problem (as expressed by the topology of a CEG) is common knowledge, and that the players are SEUM given the information available to them when they make a move. In Section 2 we introduce the semantics of the decision CEG and discuss the principle of parsimony. To illustrate how the CEG can be used for the representation and analysis of games, and also how it can be used to simplify these games, Section 3 contains a description of a 2 player game modelling a simplified version of the radicalisation scenario described above. Section 4 contains a discussion of ideas prompted by the work in earlier sections.

We have focussed here on a two person adversarial game, but note that the techniques described can be extended for use with multi-player games. We have also assumed here that we are supporting one of the two players, but because of the common knowledge assumption we have made, the qualitative results of the analysis are equally valid to this player's opponent or indeed some independent external observer.

## 2. Decision chain event graphs

### 2.1. Conditional independence, chain event graphs and causal hypotheses

Bayesian Networks (BNs) and Influence Diagrams express the conditional independence/Markov structure of a model through the presence/absence of edges between vertices of the graph. We say that a variable $X$ is independent of a variable $Y$ given $Z$ (written $X \amalg Y \mid Z$) if once we know the value taken by $Z$, then $Y$ gives us no further information for forecasting $X$. The structure of an ID can be used to produce fast algorithms for finding optimal decision strategies [6].

One advantage that CEGs have over BNs and IDs for asymmetric problems is that they can be used to represent context-specific conditional independence properties such as $X \amalg Y \mid (Z = z)$, which hold only for a subset of values of the conditioning variable.

The CEG is a function of a probability (or event) tree, having the same structure as a game tree, but with all non-leaf vertices being *chance* nodes and all edges representing outcomes of these chance nodes, rather than actions of a player, We introduce two partitions of the vertices of the tree:

- Vertices in the same *stage* have sets of outgoing edges representing the same collections of possible outcomes, and have the same probabilities of these outcomes.
- Vertices in the same *position* have sets of outgoing subpaths representing the same collections of possible complete futures, and have the same probabilities of these futures.

These equivalence classes encode (context-specific) conditional independences as follows: Given arrival at one of the vertices in a particular stage, the next development is independent of precisely which vertex has been arrived at. Given arrival at one of the vertices in a particular position, the complete future is independent of precisely which vertex has been arrived at.

Our CEG is then produced from the tree by combining (or coalescing) vertices which are in the same position. Vertices in the same stage are generally given the same colour, and equivalent edges emanating from vertices in the same stage are generally also given the same colour. The stages and positions between them encode the full conditional independence/Markov structure of our model. More detailed definitions are given in [15].

In [22], Pearl discusses the assumptions under which BNs can be considered causal (a more decision-theoretic approach to graphical modelling is considered in [23]). We have shown that under similar assumptions CEGs can also be considered as causal [18]. Heuristically this means that the model specified by a CEG continues to be valid when particular variables are manipulated. Such a hypothesis is a particularly natural one to entertain in decision problems, where a decision maker (DM) by choosing a specific action at some point can be thought of as manipulating a specific variable. The hypothesis is also a natural one to entertain in a game whose underlying structure is common knowledge and where each player is able to manipulate their own decisions to a