European Journal of Operational Research 239 (2014) 490–503

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Decision-network polynomials and the sensitivity of decision-support models

Emanuele Borgonovo *, Fabio Tonoli

Department of Decision Sciences, Bocconi University, Milan, Italy

ARTICLE INFO

Article history:
Received 28 June 2013
Accepted 7 May 2014
Available online 17 May 2014

Keywords:
Decision analysis
Influence diagrams
Decision trees
Bayesian networks
Sensitivity analysis

ABSTRACT

Decision makers benefit from the utilization of decision-support models in several applications. Obtaining managerial insights is essential to better inform the decision-process. This work offers an in-depth investigation into the structural properties of decision-support models. We show that the input–output mapping in influence diagrams, decision trees and decision networks is piecewise multilinear. The conditions under which sensitivity information cannot be extracted through differentiation are examined in detail. By complementing high-order derivatives with finite change sensitivity indices, we obtain a systematic approach that allows analysts to gain a wide range of managerial insights. A well-known case study in the medical sector illustrates the findings.

1. Introduction

Decision trees, influence diagrams, Bayesian networks and event trees support the solution of decision analysis problems in several applications. Their use is nowadays facilitated by a number of software programs (see Bielza, Gomez, & Shenoy, 2011; Jensen & Nielsen, 2007).[1] Computer implementation enables analysts to develop sophisticated codes that incorporate a variety of aspects of the problems under investigation (Dillon, Paté-Cornell, & Guikema, 2003). However, model complexity exposes analysts and decision-makers to the risk of a partial understanding of the model input–output response. Then, deriving insights about the structure of the model becomes essential to make robust conclusions and inferences.

Researchers have developed methods for exploring the informational content of decision-support models. For Bayesian networks several of the most recent findings rest on the fundamental result that the input-mapping in Bayesian networks is a multilinear polynomial (Castillo, Gutiérrez, & Hadi, 1996, 1997). For instance, multilinearity is crucial for arithmetic and decision circuits (Bhattacharjya & Shachter, 2012; Darwiche, 2003). However, *there are unique challenges in sensitivity analysis for influence diagrams due to the non-linearities created by maximization operations for making decisions* (Bhattacharjya & Shachter, 2010, p. 1). Indeed, this issue is transversal to all decision-support models that include a maximization (or minimization) operator. The maximization operation induces piecewise-definiteness and impairs differentiation. Thus, properties of Bayesian networks cannot be transferred directly to models such as influence diagrams, decision trees and decision circuits. The difficulties associated with differentiation then raise broader issues concerning the methodology for deriving managerial insights from decision-support models.

In this work, we conduct a systematic investigation into the mathematical properties of decision-support models. We show that the decision-theoretical principles underlying their construction (Savage, 1954) lead to a piecewise defined input–output mapping. Each piece corresponds to an available strategy and is multilinear in probabilities and utilities. We call the mapping a decision-network polynomial.

Non-differentiability occurs at those values of the model inputs (probabilities/utilities) for which the decision-maker is indifferent among alternative strategies. Using the terminology of Howard (1968) (see also Bhattacharjya & Shachter, 2010),[2] this result suggests that differentiation finds its natural application in an *open-loop* analysis, i.e., when the preferred strategy (or any other strategy) is under scrutiny. In a *closed-loop* analysis differentiation might not be possible. We then employ sensitivity measures based

* Corresponding author. Tel.: +39 0258365608; fax: +39 0258365630.
E-mail address: emanuele.borgonovo@unibocconi.it (E. Borgonovo).

[1] See also the associated website http://bndg.cs.aau.dk/ for a list of software programs.

[2] Drawing from the systems analysis literature, Howard, 1968 *refers to varying a parameter and computing the certain equivalent at a fixed strategy as open loop analysis; when the strategy is allowed to vary by re-evaluating the decision situation, it is closed loop analysis* (Bhattacharjya & Shachter, 2010, p. 3).

http://dx.doi.org/10.1016/j.ejor.2014.05.015
0377-2217/

on an orthogonal decomposition via finite difference operators (Borgonovo, 2010), that do not require differentiability. These finite change sensitivity indices allow us to obtain a clear understanding of the model response, by apportioning the change in expected utility to the individual effects and interactions among the model inputs.

The next step is the analysis of the model structure when consequences are monetary and certainty equivalents are the output of interest. Findings show that the input–output mapping remains piecewise-defined but each piece is composite multilinear. We link high order derivatives of certainty equivalents polynomials analytically to the derivatives of the corresponding decision-network polynomials. Finite change sensitivity indices acquire a direct interpretation as the monetary gain (loss) associated with the model input variations.

The well known case-study of Felli and Hazen (2004) helps us illustrating how the approach can be used for: (i) understanding direction of change, (ii) quantifying the relevance of interactions; and (iii) identifying the model inputs *on which to focus managerial attention during implementation* (Eschenbach, 1992, pp. 40-41).

The remainder of the paper is organized as follows. Section 2 offers a literature review. Section 3 presents decision-network polynomials. Section 4 investigates the link between indifference and differentiation. Section 5 discusses the presence of interactions and specializes finite change sensitivity indices to the case of decision-network polynomials. Section 6 discusses results for certainty equivalents. Section 7 illustrates the derivation of decision-making insights through a case study in the medical sector. Section 8 offers conclusions. All proofs are in Appendix A.

2. Review and taxonomy

This section offers a concise overview of relevant literature on decision support models and their sensitivity analysis. For a broad overview of these models we refer to Bielza et al. (2011).

Bayesian networks are among the most widely used decision-support models for the factorization of probability distributions. Their applications range from reliability analysis to genetics (Darwiche, 2010; Jensen & Nielsen, 2007). Technically, a Bayesian network is a directed acyclic graph $G = (N, A)$, where N and A are the sets of associated nodes and arcs. N contains chance nodes. Each chance node represents a random variable. A contains arcs or edges, which join pairs of nodes. The lack of an arc signifies probabilistic independence, while the presence of an arc signifies a possible probabilistic dependence. A conditional probability table (CPT) is assigned to each random variable. Algorithms for evaluating Bayesian networks have been widely studied. We refer to Jensen and Nielsen (2007) and Darwiche (2010) for thorough overviews.

A second class of decision-support models for the factorization of probability distributions are event trees. Event trees are often applied in reliability analysis in conjunctions with fault trees. Papazoglou (1998) offers a rigorous mathematical formalization. Marsh and Bearfiled (2008) show that there is always a unique Bayesian network corresponding to an event tree, but there might be multiple event trees corresponding to a Bayesian network. The event tree is unique once the Markovian assumptions of a Bayesian network are satisfied (Darwiche, 2010).

Influence diagrams are Bayesian networks *augmented with decision-nodes and value nodes, where value nodes have no descendant* (Nielsen & Jensen, 2003, p. 223). Thus, the set of nodes is now partitioned into *value, chance and decision nodes* (Shachter, 1986, p. 874). Value (or utility) nodes conclude the diagram and display the decision-maker's utility (or payoff) over consequences. Arcs evidence the flow of information, besides probabilistic dependence Shachter (1986, p. 417). As far as chance nodes are concerned, the corresponding random variables can be discrete (Howard & Matheson, 1981; Shachter, 1986, 1988) or continuous (Cobb & Shenoy, 2008; Shachter & Kenley, 1989). Influence diagrams fully retain the probabilistic meaning of Bayesian networks so that the *various well established algorithms for Bayesian net evaluation can be used in influence diagram evaluation* (Qi & Poole, 1995, p. 501). Bhattacharjya and Shachter (2010) provide a thorough review on solution methods for influence diagrams.

The compact representation offered by influence diagrams does not allow us to appreciate the detailed combination of decisions and outcomes. To reveal them, we need to convert the influence diagram into the corresponding decision-tree. Under the single decision-maker and no-forgetting conditions (which we assume throughout this work) Howard and Matheson (1981), an influence diagram can be uniquely associated with a decision-tree. At the graphical level, a decision tree contains decision, chance nodes, branches and end nodes (or leaves). They display the combinations of outcomes and alternatives that lead to the end consequences. Any such combination is called a *scenario*. The size of decision trees grows exponentially with the number of nodes, which is one of their main limitations Bielza and Shenoy (1999).

Sequential decision-diagrams (Covaliu & Oliver, 1995), unconstrained influence diagrams (Jensen & Vomlelova, 2002), limited memory influence diagrams (Lauritzen & Nilsson, 2001), valuation networks (Shenoy, 1992), sequential valuation networks (Demirer & Shenoy, 2006) extend the family of decision-support models.

Given the complexity of decision-support models in practical applications, their sensitivity analysis plays a central role for extracting managerial insights. For Bayesian networks, we recall one-way, distance-based and differentiation based sensitivity analysis. One-way sensitivity analysis is *the simplest type of analysis* and it consists of *systematically varying one of the network's parameter probabilities while keeping all other parameters fixed (*van der Gaag, Renooij, & Coupé, 2007, p. 104*)*. The works of Castillo et al. (1996, 1997) obtain analytically the sensitivity function, namely the dependence of the posterior probability of interest on the network parameter under scrutiny. An intrinsic limitation of any one-way approach is its reliance on the variation of a single parameter. Chan and Darwiche (2005) extend the robustness question to simultaneous variations. They introduce a new metric for bounding *global belief changes that result from either the perturbation of local conditional beliefs or the accommodation of soft evidence* (Chan & Darwiche, 2005). Brosnan (2006) relies on Shannon's entropy and the Kullback–Leibler divergence for answering the same questions. As to differentiation, Chan and Darwiche (2004) discuss the determination of the individual or simultaneous changes in Bayesian network-parameters that ensure the satisfaction of a query constraint using a differential approach. In Bayesian networks, partial derivatives can also be obtained by differentiating sensitivity functions; the result is called *sensitivity value* in van der Gaag et al. (2007, p. 104). In Blackmond-Laskey (1995) partial derivatives are used to focus elicitation efforts on the most important model inputs. Darwiche (2003) and Park and Darwiche (2004), augment differentiation with probabilistic semantics. They also introduce numerically efficient ways of obtaining first and second order derivatives by upwards and downward passes in arithmetic circuits.

In influence diagrams, the same sensitivity questions can be asked, but we have two notable differences. The first distinction is marked by the need to consider *decision sensitivity* besides *value sensitivity*. That is, changing the model inputs does not only modify the value of the decision-problem, but may cause the decision-maker to change strategy. Robustness then becomes the problem of assessing the region in the input parameter space over which the optimal policy is invariant. The second distinction is marked by the presence of utilities in influence diagrams, which play the