International Journal of Approximate Reasoning ••• (••••) •••–•••

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

An improved method for solving Hybrid Influence Diagrams

Barbaros Yet[a,*], Martin Neil[b], Norman Fenton[b], Anthony Constantinou[b], Eugene Dementiev[b]

[a] Department of Industrial Engineering, Hacettepe University, Ankara, 06800, Turkey
[b] School of Electronic Engineering and Computer Science, Queen Mary University of London, London, E1 4NS, UK

ARTICLE INFO

Article history:
Received 28 June 2017
Received in revised form 25 January 2018
Accepted 29 January 2018
Available online xxxx

Keywords:
Hybrid Influence Diagrams
Bayesian Networks
Dynamic Discretization

ABSTRACT

While decision trees are a popular formal and quantitative method for determining an optimal decision from a finite set of choices, for all but very simple problems they are computationally intractable. For this reason, Influence Diagrams (IDs) have been used as a more compact and efficient alternative. However, most algorithmic solutions assume that all chance variables are discrete, whereas in practice many are continuous. For such 'Hybrid' IDs (HIDs) the current-state-of-the-art algorithms suffer from various limitations on the kinds of inference that can be performed. This paper presents a novel method that overcomes a number of these limitations. The method solves a HID by transforming it to a Hybrid Bayesian Network (HBN) and carrying out inference on this HBN using Dynamic Discretization (DD). It generates a simplified decision tree from the propagated HBN to compute and present the optimal decisions under different decision scenarios. To provide satisfactory performance the method uses 'inconsistent evidence' to model functional and structural asymmetry. By using the entire marginal probability distribution of the continuous utility and chance nodes, rather than expected values alone, our method also enhances decision analysis by offering the possibility to consider additional statistics other than expected utility, such as measures of risk. We illustrate our method by using the oil wildcatter example and its variations with continuous nodes. We also use a financial score to combine risk and return measures, for illustration.

1. Introduction

An Influence Diagram (ID) is a graphical probabilistic model that offers a general and compact representation of decision making problems under uncertainty [18,22]. Fig. 1 shows an ID of the oil wildcatter problem [39]. In this problem the wildcatter is searching for oil, and has to decide whether to *drill* (D) a particular site. He is uncertain about the *quantity of oil available* (O). The wildcatter can make a *seismic test* (T), which can reveal more information about presence of oil, but the *result of this test* (R) is not totally accurate. In this ID rectangles represent decision nodes, ellipses represent chance nodes, and diamonds represent utility nodes. Each decision node represents a decision making stage, each chance node represents a random variable, and each utility node has an associated table or a continuous probability distribution that defines the utility values based on the states of its parents. Chance nodes can either be observed or not – for example, the chance node O generally cannot be observed, whereas a chance node, such as the test result R, may be observed if the

* Corresponding author.
E-mail address: barbaros.yet@hacettepe.edu.tr (B. Yet).

https://doi.org/10.1016/j.ijar.2018.01.006
0888-613X/

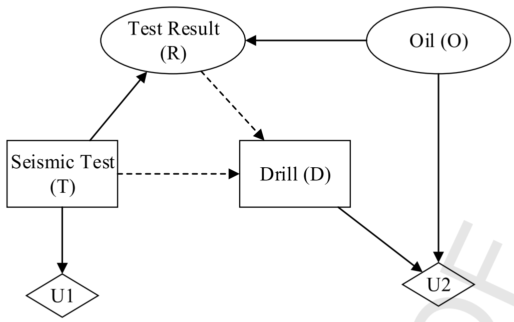

Fig. 1. Influence diagram of oil wildcatter example.

decision maker decides to go ahead with a decision (in this case to undertake the seismic test T). Incoming arcs to chance or utility nodes represent causal, deterministic or associational relations between the node and its parents. Incoming arcs to decision nodes (shown by a dashed line) are 'informational' arcs, representing the assumption that the state of any parent node must be known before the decision is made. Informational arcs also specify a partial sequential order of decisions and observations.

Generally, the chance, decision and utility nodes of an ID can be discrete or continuous variables. Many real-world problems can be represented by using a mixture of both discrete and continuous variables. An ID used to represent such a problem is called a Hybrid ID (HID), and, as we explain in Section 2, the current state-of-the-art algorithms suffer severe limitations when attempting to solve HIDs. This paper describes a novel method and algorithm, to solve HIDs, designed to overcome these limitations. Our method is based on using the Dynamic Discretization (DD) algorithm [34], which was developed to solve Bayesian Networks (BNs) containing continuous and discrete variables, to solve HIDs and to provide optimal strategies in a simplified Decision Tree (DT) that contains only decision and observable chance nodes. Unlike previous algorithms, our method provides a fully automated solution for HIDs that contain continuous chance nodes with virtually any probability distribution, including non-Gaussian types, or any conditionally deterministic function of these distributions.

The paper is structured as follows: In Section 2 we describe the state-of-the-art of DTs, IDs and their algorithms, highlighting their limitations with respect to HIDs. Sections 3 and 4 describe our novel method that adapts an existing BN DD algorithm to solve HIDs, and Section 5 presents our conclusions.

2. Decision Trees (DTs), Influence Diagrams (IDs) and Hybrid Influence Diagrams (HIDs)

In this section, we discuss the advantages and limitations of previous DT, ID and HID methods for solving decision making problems under uncertainty.

2.1. Decision Trees (DTs)

DTs have traditionally been used to choose an optimal decision from a finite set of choices, which are sometimes called policies. Typically, the value being optimized is some utility function expressed for each possible outcome of the decision. A DT represents the structure of a decision problem by modeling all possible combinations of decisions and observations, usually in the particular sequence which one would expect observations and decisions to be made. DTs are composed of three types of nodes: chance nodes, decision nodes and utility nodes. Each outgoing arc from a chance node represents an outcome and is labeled with the name and the probability of this outcome. Each outgoing arc from a decision node is labeled with a decision alternative. The DT in Fig. 2 is a representation of the wildcatter ID problem shown in Fig. 1.

While DTs are a conceptually simple and popular method for decision analysis in practice there are a number of known limitations, the main ones being:

- A DT specifies all possible sequences of observations and decisions as paths from the root node to the leaf nodes. This causes the number of state combinations to grow in size exponentially as the number of decisions and outcomes increase. This means that even simple decision problems can have infeasibly large DTs especially when there are multiple unobservable chance nodes. Domain experts may not be able to build or interpret such complex DTs effectively.
- Each path from the root to a leaf of the model represents a sequence of decisions and observations, called a decision scenario. A DT assumes 'no forgetting', i.e. at any point in the DT the decision maker knows the states of all previous nodes from the root node. The sequential order between the decision nodes and chance nodes is defined according to information available at each decision making stage. However, the order between consecutive chance nodes in a tree is usually arbitrary regardless of the conditioning and informational relationships that exist in the real world, such as those that represent causality [18]. This further increases the difficulty of understanding of complex DTs as experts often describe and interpret domain knowledge by using causal statements [17].