Artificial Intelligence 205 (2013) 30–38

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

On the complexity of solving polytree-shaped limited memory influence diagrams with binary variables

Denis Deratani Mauá *, Cassio Polpo de Campos, Marco Zaffalon

Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Galleria 2, Manno, 6928 Switzerland

ARTICLE INFO

Article history:
Received 1 March 2013
Received in revised form 10 October 2013
Accepted 18 October 2013
Available online 24 October 2013

Keywords:
Decision theory
Influence diagrams
Decision networks
Probabilistic planning
Computational complexity

ABSTRACT

Influence diagrams are intuitive and concise representations of structured decision problems. When the problem is non-Markovian, an optimal strategy can be exponentially large in the size of the diagram. We can avoid the inherent intractability by constraining the size of admissible strategies, giving rise to limited memory influence diagrams. A valuable question is then how small do strategies need to be to enable efficient optimal planning. Arguably, the smallest strategies one can conceive simply prescribe an action for each time step, without considering past decisions or observations. Previous work has shown that finding such optimal strategies even for polytree-shaped diagrams with ternary variables and a single value node is NP-hard, but the case of binary variables was left open. In this paper we address such a case, by first noting that optimal strategies can be obtained in polynomial time for polytree-shaped diagrams with binary variables and a single value node. We then show that the same problem is NP-hard if the diagram has multiple value nodes. These two results close the fixed-parameter complexity analysis of optimal strategy selection in influence diagrams parametrized by the shape of the diagram, the number of value nodes and the maximum variable cardinality.

1. Introduction

Planning with influence diagrams in partially observable domains suffers from the so-called curse of history: the size of an optimal strategy grows exponentially large with the number of look-ahead steps considered. Lauritzen and Nilsson proposed using limited-memory strategies to avoid the complexity blow up, coining the term *limited memory influence diagrams* (limids) [1]. Limids relax the *no-forgetting* requirement of traditional influence diagrams, and require (implicitly) that the maximum size of an optimal strategy be given as part of the input. In fact, the use of finite-state controllers, which in the influence diagram formalism corresponds to limited memory, has long been seen as a viable alternative in probabilistic planning for long- or infinite-horizon problems, especially when teams are considered [2–6].

In many cases, however, limiting the agent's memory resources does not make the problem tractable, and computing optimal strategies even for structurally very simple limids can be hard. We have shown in a previous work that finding an optimal strategy for polytree-shaped limids is NP-hard even if variables are ternary and there is only a single value node [7]. Also finding a strategy whose expected utility is within a factor of 2^p of the maximum expected utility, where p denotes the number of numerical parameters in the model, is NP-hard, even in polytree-shaped diagrams (with variables taking on arbitrarily many values). On the other hand, we have shown that when the variable cardinalities and the treewidth of the diagram are bounded, a fully polynomial-time approximation scheme exists [8], and that optimal strategies can often be obtained in feasible time [9]. For the remaining cases, efficiency can be achieved at the expense of provably good accuracy

* Corresponding author.
E-mail addresses: denis@idsia.ch (D.D. Mauá), cassio@idsia.ch (C.P. de Campos), zaffalon@idsia.ch (M. Zaffalon).

0004-3702/$ – see front matter
http://dx.doi.org/10.1016/j.artint.2013.10.002

Table 1
Parametrized complexity of the decision version of the optimal strategy selection problem in limids of bounded treewidth.

Topology	Number of value nodes	Maximum variable cardinality	Complexity
polytree	one	two	P
polytree	unbounded	two	NP-complete
polytree	one	three	NP-complete
polytree	one	unbounded	NP-complete
loopy	one	two	NP-complete
loopy	unbounded	bounded	NP-complete

by local search methods [1,4] and message-passing algorithms [10,11], while branch-and-bound procedures can be used for guaranteed accuracy [12,13]. Table 1 summarizes the fixed-parameter computational complexity of the decision version of finding optimal strategies in limids of bounded treewidth with respect to topology, number of value nodes and maximum variable cardinality.[1] The first two rows contain novel results presented here. The results in the remaining rows appeared in Refs. [7] and [8].

De Campos and Ji showed that the problem of selecting an optimal strategy in limids with a single value node can be reduced to that of computing a (tight) marginal probability bound in a credal network, and subsequently addressed by credal network inference algorithms [12]. The reduction preserves the topology of the graph and the cardinality of the variables; in particular, the reduction maps polytree-shaped diagrams with binary variables into credal polytrees with binary variables [14].

Credal networks are generalizations of Bayesian networks designed to cope with set-valued specifications of the numerical parameters of the model [15]. Computing tight bounds on the marginal probability of some variable in a credal network is NP-hard even when the network is polytree-shaped and variables are ternary [16]. However, Fagiuoli and Zaffalon's 2U algorithm [17] is able to obtain tight marginal bounds in polynomial time in polytree-shaped credal networks with binary variables. Therefore, we can find optimal strategies in polytree-shaped limids with binary chance variables and a single value node by running 2U in the corresponding credal network obtained by de Campos and Ji's reduction. This immediate and important result, which we discuss in Section 3, has apparently not been noticed before in the literature.

When the diagram has multiple value nodes, the transformation of de Campos and Ji reduces the selection of an optimal strategy to the computation of a tight bound for a *sum* of marginals of variables in a credal network, which cannot be solved in polynomial time by standard inference algorithms for credal networks and more specifically by 2U. A possible workaround is to apply a preprocessing step that transforms any diagram into an equivalent diagram that contains a single value node. However, known techniques to map a limid with multiple value nodes into an equivalent limid containing a single value node do not preserve the topology [8]. In particular, they map polytrees into multiply connected graphs. Thus, while the mapping from limids into credal networks enlarges the available toolset of algorithms for solving limids with that of credal network inference, it does not help us in answering the question of whether polytree-shaped diagrams with binary variables and multiple value nodes are in fact NP-hard to solve.

In Section 4, we answer that question affirmatively by a reduction from the partition problem. This result closes the fixed-parameter complexity analysis of the optimal strategy selection problem in limids in what regards the topology of the diagram, the number of values nodes, and the cardinality of the variables: polytree-shaped limids with binary variables and a single value node can be solved in polynomial time, and relaxing any of these assumptions leads to NP-hard problems.

2. Limited memory influence diagrams

To help introduce the concepts and motivate the use of limited memory influence diagrams, consider the following illustrative example of a decision problem under uncertainty.

Consider a system whose failure depends on statistically independent events $e_1, \dots, e_n$, which occur with known probabilities $p_1, \dots, p_n$, respectively. Let E_i $(i = 1, \dots, n)$ be a binary variable denoting whether e_i occurs ($E_i = 1$ if e_i occurs), and assume that the system failure is determined by a logical function $F(E_i, \dots, E_n)$ of the causing events ($F = 1$ if the system fails). Suppose that at a given time, the system administrator can intervene to prevent any event e_i from occurring at a cost v_i. Let D_i denote the decision of intervening on event e_i. The conditional probability of the event e_i given an intervention policy for the ith event is given by $P(E_i = 1|D_i = 1) = 0$ and $P(E_i = 1|D_i = 0) = p_i$. The cost of a system failure is represented by a variable C. Thus, the expected cost of a combination of intervention policies $\delta_1, \dots, \delta_n$ is given by

$$\sum_{F,E_1,\dots,E_n} E[C|F]P(F|E_1,\dots,E_n)\prod_{i=1}^{n} P(E_i|D_i = \delta_i) + \sum_{i=1}^{n} v_i\delta_i.$$

[1] The decision version of the optimal strategy selection problem is to decide, for a given a limid, whether there is a strategy whose expected utility exceeds a given rational number.