An anytime algorithm for evaluating unconstrained influence diagrams

Manuel Luque[1], Thomas D. Nielsen[2] and Finn V. Jensen[2]

[1]Dept. Inteligencia Artificial, UNED
28040 Madrid, Spain
mluque@dia.uned.es

AALBORG UNIVERSITET
[2]Department of Computer Science
Aalborg University, Aalborg, Denmark
{tdn,fvj}@cs.aau.dk

1 Introduction

- Unconstrained influence diagrams (UIDs) [1] represent decision problems in which the order of the decisions is not linear and the decision maker is interested in the best ordering as well an optimal choice for each decision.
- Due to the complexity of the problem temporal constraints can force the decision maker (DM) to act before the solution algorithm has finished, and, in particular, before an optimal policy for the first decision has been computed.
- There is a need for an anytime algorithm that computes a strategy and at any time provides a qualified recommendation for the first decisions of the problem.

2 Unconstrained influence diagrams

- Solving a UID means establishing a policy for each decision as well as a step-policy specifying the next decision given the observations made so far.
- Jensen and Vomlelova [1] describe an algorithm for solving a UID, which utilizes an auxiliary DAG, called an S-DAG, and solves it through dynamic programming by eliminating the variables in reverse temporal order.

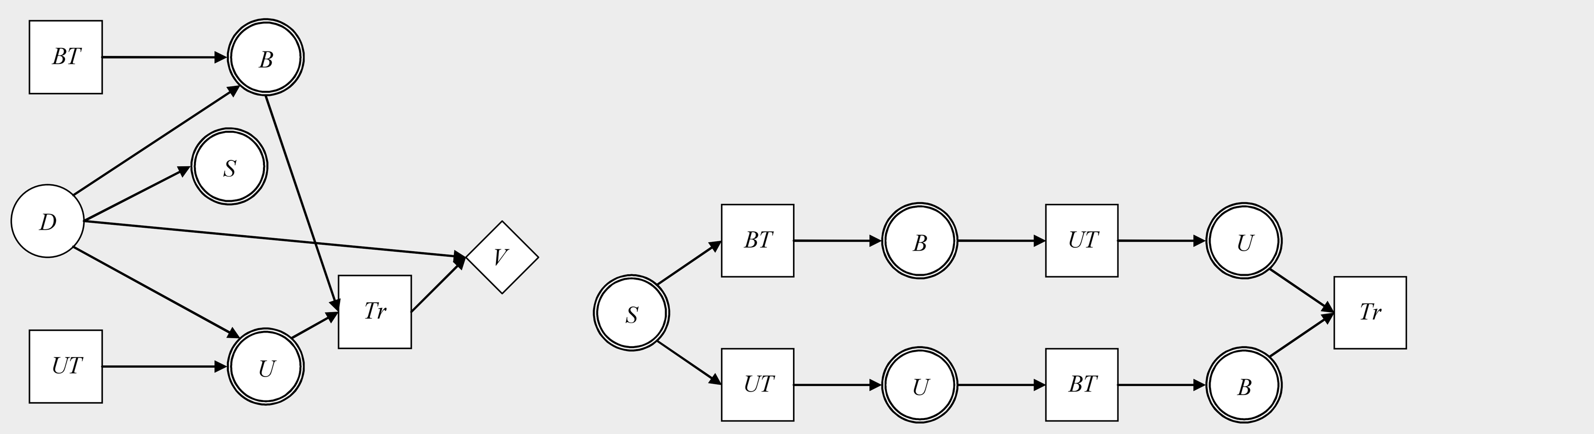

Figure 1: UID for the diabetes diagnosis problem and an S-DAG for solving it.

3 An anytime algorithm

- We propose an anytime algorithm for solving UIDs that performs a forward search in a decision tree (DT) [2] representation of the UID, guided by a heuristic function h, and building the DT from the root toward the leaves.
- h estimates an optimal policy for the decision nodes in the explored part of the DT, which form a *partial strategy*; a full strategy (*uniform extension*) can be obtained by assigning random policies to the unexplored decision nodes.

4 Selecting a heuristic function

- An upper bound h_U on the expected utility (EU) is the maximum of the utilities in the subtree rooted by X:

$$h_U(X) = \max_{l \in \mathcal{L}} \psi(\mathrm{path}(X, l)),$$

- A lower bound h_L is the EU of the uniform extension of the partial strategy.
- Nonadmissible heuristic h as a weighted linear combination of h_L and h_U:

$$h(X) = w_L(X)h_L(X) + w_U(X)h_U(X),$$

where $w_L(X) = \alpha \cdot k_X \cdot c(X)$ and $w_U(X) = \alpha \cdot d(X)$; here $c(X)$ and $d(X)$ are the number of chance and decision nodes in $\mathrm{future}(X)$, respectively.
- k_X is updated automatically as the tree is expanded.

5 Experiments

- For comparison we used the dynamic programming algorithm (DP) [1], and to test the performance of the algorithms we generated a collection of 650 random UIDs.
- The performance of the algorithm is evaluated according to the following two characteristics:
  - The frequency with which the anytime algorithm returns the correct decision options (relative to the optimal strategy) for all decision nodes down to the ith level in the decision tree.
  - The expected utility of following the strategy prescribed by the anytime algorithm or DP for the first i levels of decisions, followed by the optimal strategy for the remaining decisions.
- Time is specified relative to the time required for DP to finish.
- The reported values are normalized with the uniform strategy as baseline value, by attaining the values 0 and 1 to the uniform strategy and the optimal strategy, respectively.
- The results obtained by letting the anytime algorithm run for e.g. 50% of the time required by DP are listed in the second column in Table 1.
- From the results we clearly see that the algorithm improves over time w.r.t. all the recorded characteristics.

	25 %	**50 %**	**75 %**
EU of prescribed strategy down to lvl. 1	0,442	0,514	0,538
EU of prescribed strategy down to lvl. 2	0,609	0,769	0,865
EU of prescribed strategy down to lvl. 3	0,546	0,703	0,794
Normalized frequency of selecting optimally down to lvl. 1	0,383	0,484	0,505
Normalized frequency of selecting optimally down to lvl. 2	0,396	0,503	0,563
Normalized frequency of selecting optimally down to lvl. 3	0,291	0,381	0,428

Table 1: Results for the anytime algorithm.

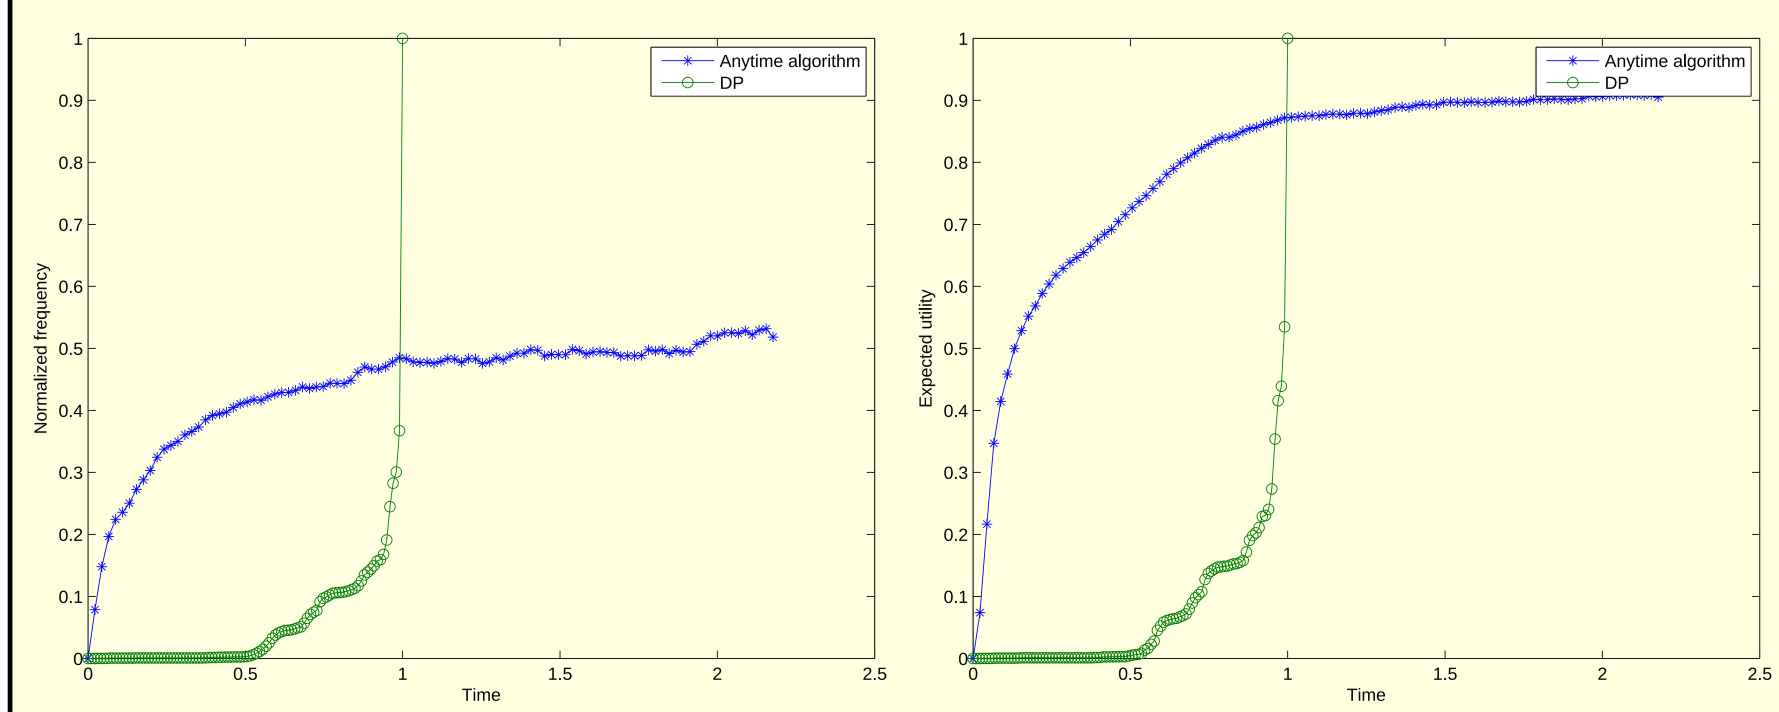

Figure 2: Comparisons of the normalized frequency of selecting optimally down to level 3 and EU of prescribed strategy down to level 3, respectively, between DP and the proposed algorithm.

6 Acknowledgements

- The first author was supported by the Department of Education of Madrid, the European Social Fund and the Spanish Ministry of Education and Science (grant TIN-2006-11152). We would like to thank Marta Vomlelova for giving us access to her UID implementation.

[1] F. V. Jensen and M. Vomlelova. Unconstrained influence diagrams. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI'02)*, pages 234–241, San Francisco, CA, 2002. Morgan Kaufmann.

[2] H. Raiffa and R. Schlaifer. *Applied Statistical Decision Theory*. MIT press, Cambridge, 1961.

The fourth European Workshop on Probabilistic Graphical Models (PGM'08).