



Temporal Relevance in Dynamic Decision Networks with Sparse Evidence*

AHMED Y. TAWFIK

School of Computer Science, University of Windsor, Windsor, Ontario, Canada

atawfik@uwindsor.ca

SHAKIL KHAN

Department of Computer Science, York University, Toronto, Ontario, Canada

skhan@cs.yorku.ca

Abstract. Dynamic decision networks have been used in many applications and they are particularly suited for monitoring applications. However, the networks tend to grow very large resulting in significant performance degradation. In this paper, we study the degeneration of relevance of uncertain temporal information and propose an analytical upper bound for the relevance time of information in a restricted class of dynamic decision networks with sparse evidence. An empirical generalization of this analytical result is presented along with a series of experimental results to verify the performance of the empirical upper bound. By discarding irrelevant and weakly relevant evidence, the performance of the network is significantly improved.

Keywords: temporal reasoning, dynamic decision networks, temporal relevance

1. Introduction

Probability theory has historically been used to model uncertainty and to make useful inference under uncertainty in various application domains. Belief networks (also known as Bayesian networks) have successfully exploited probabilistic independence to reduce the complexity of representing intricate domains [13]. To represent variables that change over time, it is possible to use a time-sliced network such that each time-slice corresponds to a time point. These time-sliced networks are known as Dynamic belief networks (DBNs) [3]. For each time in which the values of variables may change, a new slice is created. Each slice consists of a set of nodes representing values at a specific point in time. A node may have incident edges originating at other nodes within the same slice or in an earlier slice to represent the fact that a variable's value may depend on concurrent values of other variables (contemporary

influences) and on earlier values of the same and other variables (temporal or latent influences). In monitoring applications, DBNs are used to estimate the probability of hidden variables (or states) given some information about observable states as evidence. In other applications, evidence propagation across time allows past evidence to be revised in hindsight as new evidence becomes available. Exact algorithms for classes of monitoring DBNs are known. For example, the roll-up algorithm [14] works for DBN architectures that propagate effects between two consecutive time-slices in each step and dHugin algorithm [9], which uses a temporal window and smoothing operators. These algorithms have enabled the use of DBNs in a variety of applications including intelligent tutoring systems [12], planning [4], and clinical decision making [10].

Unfortunately, belief network inference is NP-hard in the worst case [1]. Due to the computational complexity of reasoning with dynamic decision networks, the performance of an intelligent system that uses these graphical structures deteriorates as the amount of information available increases. Knowing how an

*Work supported in part through a Discovery Grant from the Natural Sciences and Engineering Research Council, Canada.

additional piece of evidence may affect a decision or a query would determine whether it is worth including [8]. A piece of evidence is more valuable if it is likely to change a decision. However, information that does not change any decision may still be of some value if its effect on utilities is not negligible. The present treatment extends the previous results on the degeneration of relevance of temporal information over time [15, 16] to dynamic decision networks with sparse evidence.

Section 2 introduces the fundamental notion of probabilistic relevance of information as it applies to DDNs. An analytical upper bound on relevance in a simple DDN is presented in Section 3. Section 4 discusses information relevance to decision making as opposed to the effect of information on utilities. Section 5 proposes an empirical generalization for the analytical result in Section 3. In Section 6, experimental validation of the empirical generalization is conducted followed by some conclusions in Section 7.

2. Relevant and Irrelevant Information

According to [16] relevant information, Θ_Q of a set of assertions Θ , and a query Q (possibly a set of queries) can be defined as the minimal subset of Θ , such that the query Q follows from Θ_Q , if it follows from Θ . Moreover, a probabilistic definition of irrelevance requires that for all outcomes of the query Q , the probability of the query given Θ , and Θ_Q remains the same. In other words,

$$P(Q | \Theta) = P(Q | \Theta_Q).$$

This idea can be extended to utilities. A decision theoretic definition of irrelevance requires that for all possible choices of the action A , the utility of the action given Θ , (i.e. $U(A | \Theta)$) and $U(A | \Theta_Q)$ remains the same. In other words,

$$U(A | \Theta) = U(A | \Theta_Q).$$

According to the commonsense law of inertia [11] a state persists indefinitely. But this idealization is not really useful in practice, particularly under uncertainty. In a time-sliced decision network, the relevance of information gradually degenerates as time evolves. Intuitively, a variable that has been observed once and is not continuously monitored may change with some probability over a given time interval. As time passes,

the uncertainty builds up and the relevance of the observation degenerates. Hence, a time-sliced decision network, that has a single evidence node, can be divided into two parts (periods) namely a relevance period and an irrelevance period. The evidence affects the decision during its relevance period only.

Independence captures a clear sense of mutual irrelevance. To identify a class of irrelevance that captures the relevance degeneration with time due to the uncertain dynamic nature of change, a weaker relevance criterion is necessary. If the maximum change that an assertion θ_j in Θ , at time t_j can induce on the utility of action a_{t_i} at time t_i is less than a small value δ , then t_i and t_j are temporally extraneous with respect to a_{t_i} .

Definition 1. For a binary variable (or a conjunction of binary variables) Θ and an action a_{t_i} , the degree of relevance δ of Θ at time t_j with respect to a_{t_i} at time t_i , can be defined as the smallest δ that satisfies the inequality:

$$|U(a_{t_i} | \Theta_j) - U(a_{t_i} | \neg \Theta_j)| \leq \delta.$$

If Θ_j or $\neg \Theta_j$ contain disjunctions of mutually exclusive outcomes (e.g. the distinct values of variables in Θ_j are not binary) it is more efficient and convenient to perform pairwise comparisons of each pair of these outcomes. Hence, we extend the definition of relevance to use pairwise comparisons as follows:

Definition 2. The degree of relevance of factor Θ_j with respect to action a_{t_i} is δ iff for all possible assignments of Θ_j , the maximum change in the utility $U(a_{t_i} | \theta_j)$ is less than δ .

In the above definition, δ represents the strength of the degree of relevance. Since our interest is in weak temporal relevance with reasonably small δ values, we will give another definition that allows us to ignore very weak relevance:

Definition 3. The theory Θ can be divided into a relevant subset Θ_Q and an extraneous subset $\Theta - \Theta_Q$, the relevant subset Θ_Q assesses the utility of Q within accuracy δ iff for any conjunction (possibly singleton) the following inequality is satisfied:

$$\forall q \subset Q, U(q | \Theta) - U(q | \Theta_Q) \leq \delta.$$

In this study of irrelevance, our goal is to identify the irrelevant information and define the relevant sub-theory Θ_Q . By precisely identifying Θ_Q , it is possible to improve the performance of intelligent systems by reducing the size of the knowledgebase that must be considered before answering the query or making a decision.

3. An Upper Bound on Relevance Time in Dynamic Decision Networks

An upper bound on the time duration T for a single variable time-sliced belief network is presented in [16]. Here we proceed with a similar analysis to show that there exists a time duration T such that the utility of an action a at time $t > t_0 + T$ changes by at most δ depending on the evidence available at time t_0 . The analysis presented here, unlike in previous work, takes into account the effect of utilities and how the utilities affect decisions in a dynamic decision network.

Theorem 1. *In a time sliced decision network, consider a fluent C with states c_i and $\neg c_i$ and the four transition probabilities: $P(\neg c_{i+\Delta} | c_i) = p_1$, $P(c_{i+\Delta} | \neg c_i) = p_2$, $P(c_{i+\Delta} | c_i) = 1 - p_1$, $P(\neg c_{i+\Delta} | \neg c_i) = 1 - p_2$, such that $0 \leq p_1, p_2 \leq 1$. If the system is in state c_i then the fluent is true at time i . Let the utility of the decision D at time t be $U_t(D)$ and let the probability of the fluent is in state C_t at time t be $P(C_t)$. We claim that for any $\delta \ll 1$, there exists T such that $\forall t \geq T$, $|U_t(D | c_0) - U_t(D | \neg c_0)| \leq \delta$.*

Proof:

$$U_t(D | c_0) = U(D | c_t)P(c_t | c_0) + U(D | \neg c_t)P(\neg c_t | c_0) \quad (1)$$

$$= (U(D | c_t) - U(D | \neg c_t))P(c_t | c_0) + U(D | \neg c_t) \quad (2)$$

$$P(c_t | c_0) = P(c_t | c_{t-1})P(c_{t-1} | c_0) + P(c_t | \neg c_{t-1})P(\neg c_{t-1} | c_0) \quad (3)$$

$$= (1 - p_1)P(c_{t-1} | c_0) + p_2P(\neg c_{t-1} | c_0) \quad (4)$$

$$= (1 - p_1)P(c_{t-1} | c_0) + p_2(1 - P(c_{t-1} | c_0)) \quad (5)$$

$$= (1 - p_1 - p_2)P(c_{t-1} | c_0) + p_2 \quad (6)$$

Using the iteration method, we solve this recurrence relation and get,

$$P(c_t | c_0) = (1 - p_1 - p_2)^{t-1}P(c_1 | c_0) + p_2(1 - p_1 - p_2)^{t-2} + \dots + p_2(1 - p_1 - p_2) + p_2 \quad (7)$$

Summing the geometric series in (7), we have:

$$P(c_t | c_0) = p_1(1 - p_1 - p_2)^{t-1} \left[\frac{1}{p_1 + p_2} - 1 \right] + \frac{p_2}{p_1 + p_2} \quad (8)$$

Similarly,

$$P(c_t | \neg c_0) = (1 - p_1 - p_2)^{t-1} \left[1 - \frac{1}{p_1 + p_2} \right] + \frac{p_2}{p_1 + p_2} \quad (9)$$

Substituting from (8) in (2), we have

$$U_t(D | c_0) = (U(D | c_t) - U(D | \neg c_t)) \times \left(p_1(1 - p_1 - p_2)^{t-1} \left[\frac{1}{p_1 + p_2} - 1 \right] + \frac{p_2}{p_1 + p_2} \right) + U(D | \neg c_t) \quad (10)$$

Similarly,

$$U_t(D | \neg c_0) = (U(D | c_t) - U(D | \neg c_t)) \times \left(p_2(1 - p_1 - p_2)^{t-1} \left[1 - \frac{1}{p_1 + p_2} \right] + \frac{p_2}{p_1 + p_2} \right) + U(D | \neg c_t) \quad (11)$$

Subtracting (11) from (10) and taking the absolute value of the difference, we get

$$|U_t(D | c_0) - U_t(D | \neg c_0)| = |U(D | c_t) - U(D | \neg c_t)| |1 - p_1 - p_2|^t \quad (12)$$

Hence,

$$\delta \leq |U(D | c_t) - U(D | \neg c_t)| |1 - p_1 - p_2|^T \quad (13)$$

from the above inequality,

$$T \leq \frac{-\log \delta + \log |U(D | c_t) - U(D | \neg c_t)|}{-\log |1 - p_1 - p_2|} \quad (14)$$

□




Figure 1. A simple decision network with observable nodes ‘Sense’, hidden nodes ‘State’ and action nodes ‘Action’.

Intuitively, this theorem implies that information at a particular time may not help in taking a decision at another later time if the period between the two is long enough. Here, $|U_t(D | c_0) - U_t(D | \neg c_0)|$ measures the effect of knowing C at time t_0 on the utility of D at time t . In other words, having $U_t(D | c_0) \approx U_{t-1}(D | c_0)$ implies that $U_t(D | c_0)$ has converged.

4. Convergence of Decisions

So far we have considered utility convergence. Here we turn our attention to decision convergence. In a time-sliced decision network, our goal at time t is to pick a decision D that maximizes the utility U_t . Here, the question is when does D become extraneous of c_0 (a set of evidence variables at time t_0)? In other words, for a one variable binary system, what is the value of T' so that we can pick a D such that

$$\forall t > T', \quad U_t(D | c_0)P(c_0) \approx U_t(D | \neg c_0)P(\neg c_0),$$

and

$$U_t(\neg D | c_0)P(c_0) \approx U_t(\neg D | \neg c_0)P(\neg c_0)?$$

We claim that decisions converge faster than utilities in most cases. In this section, we give some arguments supporting this claim. In Section 4, we present some experimental results comparing decision and utility convergence. Consider the dynamic decision network in Fig. 1. Suppose that in the above network, all the variables have two states and we have two choices of decisions namely D and $\neg D$. Our goal is to find a time duration T' such that the decision is not changed by evidence at time t_0 . In the present discussion we assume stationarity of the probability distribution as well as the rewards used to calculate the utilities of the value nodes. Stationarity is a common assumption in Markov decision processes. This assumption implies that the probability distribution and rewards do not change with time.

Consequently, if we only have evidence at t_0 , the utilities will be either monotonically increasing or monotonically decreasing with time according to the qualitative probability model in [17]. Given that the utilities $U_t(D | c_0)$ and $U_t(D | \neg c_0)$ converge and so do the utilities $U_t(\neg D | c_0)$ and $U_t(\neg D | \neg c_0)$, we have all four utilities monotonically evolving over time. They can all be evolving in the same direction (e.g. all increasing) or they may evolve in different directions (e.g. three increasing and one decreasing). Here, we examine three cases to illustrate the convergence behavior of decisions with respect to that of utilities. It is obvious that whenever utilities converge, the decision also converges because typically, a rational agent selects the decision that maximizes the expected utility. The cases are selected from a set of randomly generated simple networks like the one shown in Fig. 1. The first two cases show that decisions converge faster than utilities. In the first case the decision converges relatively faster than in the second case because the effect of the initial condition at time t_0 is more persistent. In the third case decision and utilities converge at approximately the same time. It is important to note that it is possible to construct a case where the optimal decision changes after the utilities converge. For example, the utilities $U_t(D | c_0)$ and $U_t(\neg D)$ may intersect after their convergence. However, this change does not depend on the earlier observations at t_0 .

Case 1. First assume that $U_t(D | c_0)$ and $U_t(D | \neg c_0)$ are both monotonically increasing while $U_t(\neg D | c_0)$ and $U_t(\neg D | \neg c_0)$ are monotonically decreasing. This situation is depicted in Fig. 2.

Here we can see that decision converges between time slices 8 and 9; also notice that the utility converges between time slices 18 and 19. So clearly, decision converges faster than utility in this case.

Case 2. Consider the case when $U_t(D | c_0)$ and $U_t(\neg D | c_0)$ are increasing while $U_t(D | \neg c_0)$ and

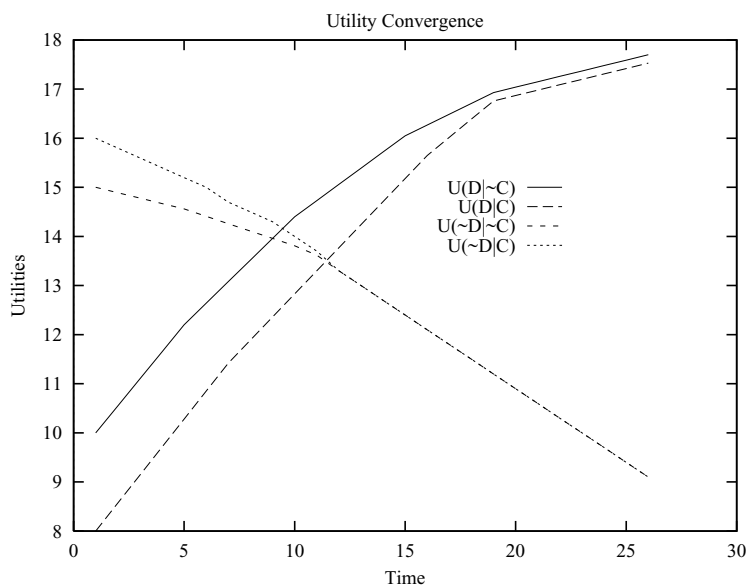


Figure 2. Monotonically increasing and decreasing utilities.

$U_t(D | c_0)$ are decreasing. This situation can be depicted as in Fig. 3.

In this case, the decisions are converging approximately before the utilities converge in the figure. Hence, this case also supports the claim that decisions converge faster than utilities.

Case 3. We can see from the two examples above that in many cases, decision converges faster than utility. However, there are cases where decision and utility both take the same time to converge. An example of this is presented in Fig. 4.

Here, $U_t(\sim D | c_0)$ and $U_t(\sim D | \sim c_0)$ converge very quickly. On the other hand, $U_t(D | c_0)$ and $U_t(D | \sim c_0)$ take a long time to converge. By looking at the graph in Fig. 4 we can see that in this case utility and decision both are converging almost at the same time.

To summarize, most of the time, decision converges faster than utility, but not necessarily. Section 6 presents some test results to verify the validity of an empirical generalization of the theorem in Section 3 to general dynamic decision networks.

5. Empirical Generalization

The analytical bound on information relevance in dynamic decision networks derived in Section 3 describes the behavior of a single variable system. Applying the

analytical approach used in Section 3 to generalized networks results in complex recurrence relations that do not lend themselves to analytical solutions. To generalize the result in Section 3 we adopt an empirical approach. The present work is somewhat related to the study of the cutoff phenomenon, in Markov chains. This phenomenon characterized by the rather sudden and fast convergence of some chains after a certain time, continues to be an area of mathematical research since it was introduced in [5]. Mathematicians have developed bounds on the convergence of random walks, diffusion models, card shuffles. The present study adopts a more heuristic, computationally efficient, and empirical approach to the problem.

To find a general upper limit on relevance time, let us consider the limit for the binary case in (14) presented in Section 3:

$$T \leq \frac{-\log \delta + \log |U(D | c_t) - U(D | \sim c_t)|}{-\log |1 - p_1 - p_2|}$$

This bound depend on the chosen accuracy level δ , the values of the rewards $U_t(D | c_t)$ and $U_t(D | \sim c_t)$ as well as the transition probabilities. It is expected that any sensible generalization should depend on all three entities. Moreover, we require that the proposed generalization supports the results in Section 3, as a special case. Keeping in mind that we are looking for an upper bound, we chose to use the following expression




Figure 3. Decision converges faster than utilities.

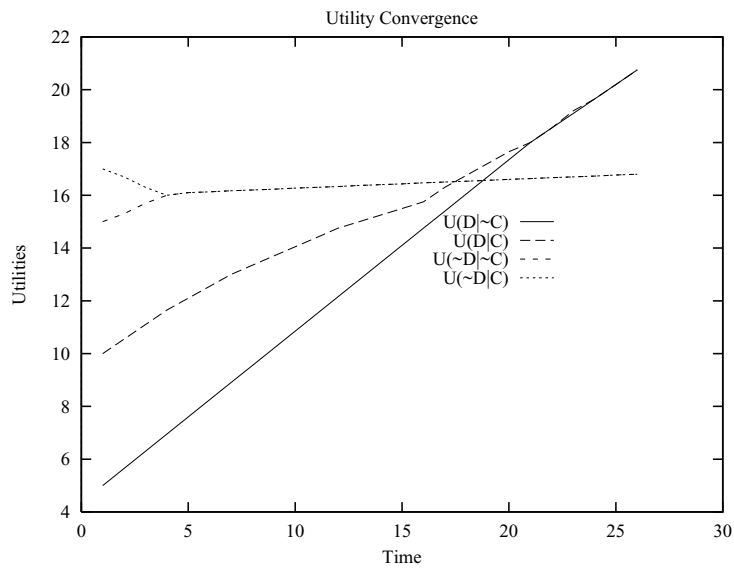


Figure 4. Utilities and decision converge at the same time.

for the empirical time bound:

$$T \leq \frac{-\log \delta + \log \text{Max} \Delta U}{-\log \text{Max} \Delta P}.$$

The notation $\text{Max}(\Delta U)$ represents the maximum difference in the utility table and $\text{Max}(\Delta P)$ represents the maximum difference in any conditional probability table associated with a temporal edge. Note that for δ

is numerically small and $\log \delta$ is negative, and so is $\log \text{Max} \Delta P$. To justify the choice of this expression, let's consider the following special cases:

- If the utilities are almost equal, $\text{Max} \Delta U = 1$, the time to convergence depends exclusively on the behavior of the probability distribution and we get the same expression as [15]. Note that the above expression assumes that $\text{Max} \Delta U$ is at least 1.

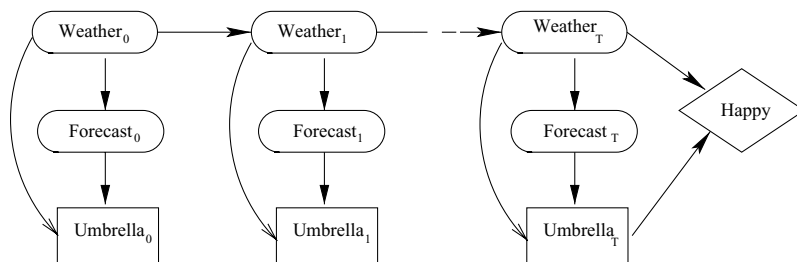


Figure 5. The weather network.

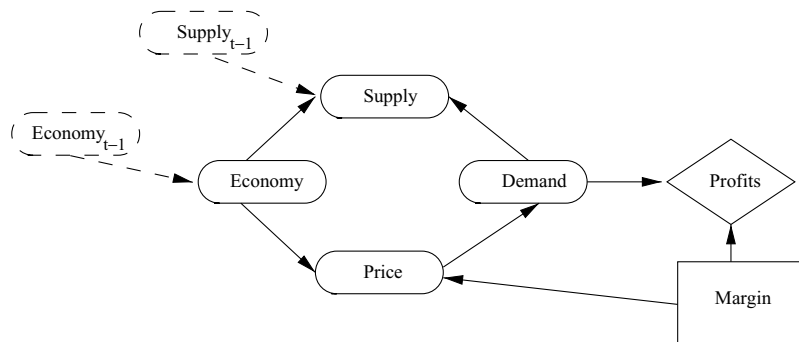


Figure 6. Car sales network.

- If the system is deterministic, we get infinite predictability as the denominator becomes 0 and the bound on T becomes infinite.
- if the probabilities are almost equal, the denominator grows large reducing the convergence time.

To assess the quality of this measure, we have conducted extensive experimental evaluation and the results obtained are promising.

6. Experimental Evaluation

We have conducted several sets of experiments using five dynamic decision network models, namely Generic, Weather, Car Sales, POMDP and Time Critical. The generic network shown in Fig. 1 has been used to study the behavior of multivalued networks. In these networks each node has more than two values. The weather network shown in Fig. 6 represents a network that is slightly more complex than the generic network as the decision to take an *umbrella* depends on two nodes *forecast* and *weather*. The utility node *happy* depends on the weather and the decision. The

car sales network extends the network in [2] by including a decision node representing the *profit margin* and a utility node representing the net *profits*. This network has four chance nodes in each time slice representing *price*, *supply*, *demand*, and *state of the economy*. Temporal edges connect supply nodes and the state of economy nodes. Figure 7 shows a time slice of this network.

The POMDP network represents a partially observable Markov process with an observable chance node (*Obs*) and an unobservable state (*State*) that determines the utility node (*Reward*) along with the decision node *Action*. Figure 7 depicts a time slice of this network. *Action* at time $t - 1$ affects *state* and *obs* at time t . Another temporal edge representing state persistence links $state_{t-1}$ to $state_t$. The network in Fig. 8 represents a time critical medical intervention in trauma cases [6]. The action node represents the possible interventions, and the evidence nodes (E1 to E5) reflect symptoms that may support or refute a diagnostic *hypothesis*. Temporal edges linking the hypothesis node in consecutive time slices capture the evolution of the case over time in response to medical intervention.




Figure 7. POMDP network.

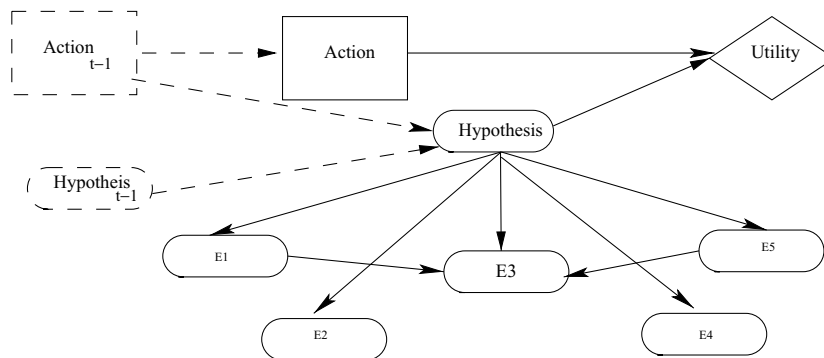


Figure 8. Time-critical intervention network.

6.1. Utility Convergence Results

In each set of experiments, we use several thousands of randomly generated dynamic decision networks as indicated in Tables 1 and 2. These test networks use randomly generated conditional probabilities and utilities. The random values are checked to eliminate inconsistent networks including networks that have utilities that differ by values smaller than the chosen δ .

All possible observations are set as evidence at time zero, one at a time, and propagated. The evaluation

algorithm keeps track of the pairwise differences in utilities due to different initial evidence. The algorithm reports the relevance time by comparing these pairwise differences to δ .

In the following discussion, the proposed predictor is considered to have under-estimated or under-predicted the relevance time if the predicted time is shorter than the actual time.

As we have mentioned previously, the complexity of decision network reasoning is in general NP-hard. So is the complexity of reasoning in DDNs. Since we use

Table 1. Performance of proposed predictor.

| Network | Nodes per slice | δ | Number of tests | Underestimated | Percentage |
|---------------|-----------------|----------|-----------------|----------------|------------|
| Generic | 2 | 0.001 | 5500 | 0 | 0 |
| Weather | 2 | 0.001 | 5129 | 47 | 0.92 |
| POMDP | 2 | 0.1 | 2718 | 0 | 0 |
| Time critical | 6 | 0.1 | 1705 | 0 | 0 |
| Car sales | 4 | 0.001 | 5161 | 0 | 0 |

Table 2. Relevance time for multivalued networks.

| Network | Number of States | δ | Iterations | Under-Estimated | Percentage |
|---------|------------------|----------|------------|-----------------|------------|
| Generic | 3 | 0.001 | 1124 | 12 | 1.07 |
| Generic | 3 | 0.01 | 2551 | 65 | 2.55 |
| Generic | 3 | 0.1 | 5677 | 174 | 3.06 |

a straightforward implementation of DDNs, we have to limit our experiments in the following manner. When calculating the actual bound, we only allow a DDN to have up to S time slices, where S is a variable that depends on the nature of the network model under consideration. The more complex the network, the smaller S gets. DDNs that take more than S time slices to converge have been eliminated from the evaluation after verifying that the predicted convergence time also exceeds S . This turned out to be true in all the eliminated test cases.

For the first set of experiments, we use the Generic network in Fig. 1 with binary valued nodes and randomly generated probabilities to validate our software used in the experimental validation. Since the proof in Section 3 provides an analytical upper bound on convergence, we have verified the implementation by running thousands of experiments using randomly generated probabilities and utilities. As expected, none of

these tests produced a case where our prediction is lower than the actual time steps needed for the utility to converge. However, most practical applications use more complex multi-node networks. The next four sets of experiments consider more realistic networks. It is clear from Table 1 that the proposed bound performed well in these test cases. More than several thousands randomly generated Time Critical, POMDP and Car Sales networks failed to produce a single case that exceeds the empirical upper bound. Moreover, the formula has under-predicted rarely for the Weather Forecast network.

Figure 9 shows the performance of the predicted bound for the Weather Forecast network (when $\delta = 0.001$ and the variables are binary). For the last three sets of experiments, we used the Generic network, but this time instead of being strictly binary valued, the variables that we used are multivalued. Unfortunately, when more values are allowed, the general formula under-predicts in some cases. However, under-prediction occurred for less than 4% of all tests. Moreover, in the cases where the bound is exceeded, it was exceeded by a small number of steps. Hence, the formula might still be useful in situations where absolute certainty is not required. Since almost all the under-predictions were off by 7 or less time slices, a ‘slush factor’ can be added to the result of the formula to guarantee zero underestimation.




Figure 9. Predicted versus actual convergence time.

Table 3. Decision convergence versus utility convergence.

| | Decision | Utility |
|-------------------------------------|----------|---------|
| Maximum time slices for convergence | 12 | 17 |
| Minimum time slices for convergence | 1 | 2 |
| Average | 1.68 | 4.67 |
| Median | 2 | 4 |

6.2. Decision versus Utility Relevance

To verify the intuitions concerning relevance of information for decision making presented earlier, the experimental evaluation results are reexamined here to test our claims. The results of this reexamination are in Table 3. As expected, in most cases, decision converges faster than utilities. Here we used the Generic network with binary valued nodes.

From Table 3, we can see that the average decision convergence is much faster than the average utility convergence. Moreover, there are some cases where decision and utility takes the same time to converge.

7. Conclusion

In dynamic decision networks with sparse evidence, the relevance of information has been shown to degenerate over time. Irrespective of available early evidence, utilities eventually converge to the same values. Consequently, earlier information becomes irrelevant to decisions that aim at maximizing utilities. Generally, irrelevance with respect to decisions occurs faster than utility convergence because the ranking of decisions is less sensitive to changes in utilities.

An analytical bound on the duration of information relevance in dynamic decision network introduced in this work, has been generalized empirically. Initial experimental results show that the proposed generalization works well for dynamic decision network with sparse evidence. This bound should allow us to ignore weakly relevant evidence to improve the computational performance without compromising the quality of the decision.

References

- G.F. Cooper, "The computational complexity of probabilistic reasoning using Bayesian belief networks," *Artificial Intelligence*, vol. 42, nos. 2/3, pp. 393–405, 1990.
- P. Dagum, A. Galper, and E. Horvitz, "Dynamic network models for forecasting," in *Proceedings of the Eighth Workshop on Uncertainty in Artificial Intelligence*, pp. 41–48, 1992.
- T. Dean and K. Kanazawa, "A model for reasoning about persistence and causation," *Computational Intelligence*, vol. 5, no. 3, 1989.
- T. Dean and M. Welman, *Planning and Control*, Morgan Kaufmann Publishers, 1991.
- P. Diaconis and M. Shahshahani, "Time to reach stationarity in the Bernoulli Laplace diffusion model," *SIAM Jour. Math. Anal.*, vol. 18, pp. 208–218, 1987.
- E. Horvitz and A. Seiver, "Time-critical action: Representation and application," in *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI-97)*, 1997.
- T. Huang, D. Koller, J. Malik, G. Ogasawara, B. Rao, S. Russell, and J. Weber, "Automated symbolic traffic scene analysis using belief networks," in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pp. 966–972, 1994.
- F. Jensen, and J. Liang, "drHugin: A system for value of information in Bayesian networks," in *Proceeding of the Fifth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU)*, Paris, France, 1994 pp. 178–183.
- U. Kjaerulff, "A computational scheme for reasoning in dynamic probabilistic networks," in *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*, pp. 121–129, 1992.
- P. Magni, "A new approach to optimal dynamic therapy planning," in *Transforming Health Care Through Informatics: Proceedings of the 1999 AMIA Annual Symposium*, 1999.
- J. McCarthy and P. Hayes, "Some philosophical problems from the standpoint of artificial intelligence," in *Machine Intelligence 4*, edited by B. Meltzer and D. Michie. Edinburgh University Press: Edinburgh, Scotland, pp. 463–502, 1969.
- R.C. Murray and K. VanLehn, "DT tutor: A decision-theoretic, dynamic approach for optimal selection of tutorial actions," in *Intelligent Tutoring Systems, 5th International Conference (ITS 2000)* edited by G. Gauthier, C. Frasson, and K. VanLehn, Springer: New York, pp. 153–162, 2000.
- J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, 1988.
- S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach* Prentice-Hall, Englewood Cliffs: New Jersey, 1995.
- A.Y. Tawfik and T. Barrie, "The degeneration of relevance in uncertain temporal domains: An empirical study," in *Advances in Artificial Intelligence: Proceedings of the Thirteenth Biennial Canadian Conference on Artificial Intelligence, (AI'2000)*, Heidelberg LNAI, Springer-Verlag: Berlin, 1822, pp. 421–431, 2000.
- A.Y. Tawfik and E. Neufeld, "Temporal reasoning and bayesian networks," *Computational Intelligence*, vol. 16, no. 3, pp. 349–377, 2000.
- A.Y. Tawfik, "Towards temporal reasoning using qualitative probabilities," in *Proceedings of FLAIRS-02, the Fifteenth Florida International Artificial Intelligence Research Symposium*, Pensacola, FL, 2002.