Temporal Relevance in Dynamic Decision Networks with Sparse Evidence

Ahmed Y. Tawfik, and Shakil Khan

School of Computer Science, University of Windsor

Windsor, Ontario N9B 3P4, Canada

atawfik@uwindsor.ca

Abstract

In this paper, we discuss 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.

I. 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 [12]. Unfortunately, belief network inference is NP-hard in the worst case [1]. To represent variables that change over time, it is possible to use a time-sliced network such that a time-slice corresponds to 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. Nodes may be connected to nodes within the same or earlier slices 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).

Decision theory extends probability theory to guide decision-making under uncertainty. Utilities are used to provide a quantitative measure of preferences among possible world states. To decide among alternative actions, the expected utility of each alternative is calculated by taking the sum of the utilities of all possible future states of the world that follow from that alternative, weighted by the probabilities of those states occurring. Decision theory holds that a rational agent chooses the alternative that maximizes the expected utility. A belief network, which consists entirely of chance nodes, can be extended into a decision network (equivalently, an influence diagram) by adding decision and utility nodes along with appropriate arcs. A dynamic decision network (DDN) is like a DBN except that it has decision and utility nodes in addition to chance nodes. DDNs model decisions for situations in which decisions, variables or preferences can change over time. These networks have been used in a variety of applications including traffic scene analysis [7], intelligent tutoring systems [11], planning [4], and clinical decision making [9].

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 additional piece of evidence may affect a decision or a query would determine whether it is worth seeking or 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 [13], [14] to dynamic decision networks with sparse evidence.

Section II 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 III . Section IV discusses information relevance to decision making as opposed to the effect of information on utilities. Section V proposes an empirical generalization for the analytical result in Section III. In Section VI, experimental validation of the empirical generalization is conducted followed by some conclusions in Section VII.

II. Relevant and Irrelevant Information

According to [14], 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 if it follows from Θ_Q. 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,

$$P(Q \mid \theta) = P(Q \mid \theta_Q).$$

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 and Θ_Q remains the same. In other words,

$$Utility(A \mid \theta) = Utility(A \mid \theta_Q).$$

According to the commonsense law of inertia [10] a state persists indefinitely. But this idealization is not really useful in practice. In a time sliced decision network, the relevance of information gradually degenerates as time evolves. Hence time sliced decision networks can be divided into two periods namely a relevance period and an irrelevance period.

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_l at time t_i is less than a small value δ, then t_i and t_j are temporarily extraneous with respect to a_l.

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

$$\left|Util(a_{l_i} \mid \theta_j) - Util(a_{l_i} \mid \neg\theta_j)\right| \le \delta \cdot$$

If θ_j or $\neg\theta_j$ contain disjunctions of mutually exclusive outcomes such as the distinct values of variables in θ_j are not binary, it is more efficient and convenient to perform pairwise comparisons of each two 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_l is δ iff for all possible assignments of θ_j, the maximum change in the utility $Util(a_l \mid \theta_j)$ 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 Θ_E. Θ_Q answers the query Q with accuracy δ iff for any conjunction (possibly singleton),

$$q \subset Q, \left|Util(q \mid \Theta) - Util(q \mid \Theta_Q)\right| \le \delta \cdot$$

In this study of irrelevance, our goal it to identify the irrelevant information and define the relevant subtheory 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.

III. An upper bound for time sliced decision networks

An upper bound on the time duration T for a single variable time sliced belief network is presented in [14]. 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 evidences available at time t_0.

Theorem

In a time sliced decision network, consider a fluent C_i with states C_i and $\neg C_i$ and the four transition probabilities:

$$P(\neg C_{i+\Delta} \mid C_i) = p_1, \quad P(C_{i+\Delta} \mid \neg C_i) = p_2,$$
$$P(C_{i+\Delta} \mid C_i) = 1 - p_1, \text{ and } P(\neg C_{i+\Delta} \mid \neg C_i) = 1 - p_2$$

such that $0 < p_1, p_2 < 1$. 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 is $U_t(D)$ and let the probability of the fluent is in state C_t at time t be $P(C_t)$. We claim that for any $\delta << 1$, there exists T such that $\forall t \ge T, \left|U_t(D \mid C_0) - U_t(D \mid \neg C_0)\right| < \delta$.

Proof.

$$\begin{aligned} U_t(D \mid C_0) &= U(D \mid C_t)P(C_t \mid C_0) + U(D \mid \neg C_t)P(\neg C_t \mid C_0) \\ &= U(D \mid C_t) \times P(C_t \mid C_0) + U(D \mid \neg C_t)(1 - P(C_t \mid C_0)) \\ &= (U(D \mid C_t) - U(D \mid \neg C_t))P(C_t \mid C_0) + U(D \mid \neg C_t) \end{aligned}$$

Now,

$$\begin{aligned} P(C_t \mid C_0) &= P(C_t \mid C_{t-1}) \times P(C_{t-1} \mid C_0) + P(C_t \mid \neg C_{t-1}) \times P(\neg C_{t-1} \mid C_0) \\ &= (1 - p_1) \times P(C_{t-1} \mid C_0) + p_2 \times P(\neg C_{t-1} \mid C_0) \\ &= (1 - p_1) \times P(C_{t-1} \mid C_0) + p_2 \times (1 - P(C_{t-1} \mid C_0)) \\ &= (1 - p_1 - p_2) \times P(C_{t-1} \mid C_0) + p_2 \end{aligned}$$

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

$$\begin{aligned} P(C_t \mid C_0) &= (1 - p_1 - p_2)^{t-1} \times P(C_1 \mid C_0) + p_2 \times (1 - p_1 - p_2)^{t-2} + \\ &\quad p_2 \times (1 - p_1 - p_2)^{t-3} + \cdots + p_2 \times (1 - p_1 - p_2) + p_2 \end{aligned}$$

Summing the geometric series and substituting for $P(C_1 \mid C_0)$ we have,

$$P(C_t \mid C_0) = p_1 \times (1 - p_1 - p_2)^{t-1} \times \left[\frac{1}{p_1 + p_2} - 1\right] + \frac{p_2}{p_1 + p_2}$$

Similarly,

$$\begin{aligned} P(C_t \mid \neg C_0) &= (1 - p_1 - p_2)^{t-1} \times P(C_1 \mid \neg C_0) + p_2 \times (1 - p_1 - p_2)^{t-2} + \\ &\quad p_2 \times (1 - p_1 - p_2)^{t-3} + \cdots + p_2 \times (1 - p_1 - p_2) + p_2 \\ &= p_2 \times (1 - p_1 - p_2)^{t-1} \times \left[1 - \frac{1}{p_1 + p_2}\right] + \frac{p_2}{p_1 + p_2} \end{aligned}$$

By Substituting for $P(C_t \mid C_0)$, we have,

$$U_t(D \mid C_0) = (U(D \mid C_t) - U(D \mid \neg C_t)) \times \left(p_1 \times (1 - p_1 - p_2)^{t-1} \times \left[\frac{1}{p_1 + p_2} - 1\right] + \frac{p_2}{p_1 + p_2}\right) + U(D \mid \neg C_t)$$

Similarly, we have,

$$U_t(D \mid \neg C_0) = (U(D \mid C_t) - U(D \mid \neg C_t)) \times \left(p_2 \times (1 - p_1 - p_2)^{t-1} \times \left[1 - \frac{1}{p_1 + p_2}\right] + \frac{p_2}{p_1 + p_2}\right) + U(D \mid \neg C_t)$$

Substituting by the above two expressions and taking the absolute value, we have,

$$\left|U_t(D \mid C_0) - U_t(D \mid \neg C_0)\right| = \left|U(D \mid C_t) - U(D \mid \neg C_t)\right| \times \left|1 - p_1 - p_2\right|^t$$

Hence,

$$\delta \ge \left|U(D \mid C_t) - U(D \mid \neg C_t)\right| \times \left|1 - p_1 - p_2\right|^T$$

from the above equation, we have,

$$T \le \frac{-\log \delta + \log\left|U(D \mid C_t) - U(D \mid \neg C_t)\right|}{-\log\left|1 - p_1 - p_2\right|}.$$

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, $\left|U_t(D \mid C_0) - U_t(D \mid \neg C_0)\right|$ measures the effect of C_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.

IV. 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 that 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, $\forall t \geq T'$,

$$U_t(D|C_0)P(C_0)+U_t(D|\neg C_0)P(\neg C_0)>U_t(\neg D|C_0)P(C_0)+U_t(\neg D|\neg C_0)P(\neg C_0)$$

We claim that decision converges faster than utilities in most cases. In this Section, we give some arguments supporting this claim. In Section VI, we present some experimental results comparing decision and utility convergence.

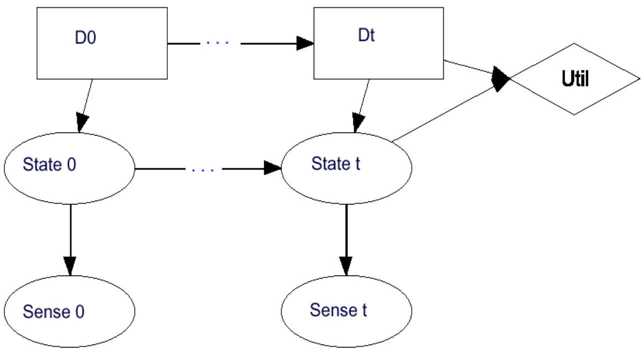

Figure1. A simple decision network

Consider the dynamic decision network in Figure 1. Suppose that in the above network, all the variables have two states and we have two choices of decisions namely *D* and ¬*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 utility 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 [15]. 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 direction (e.g. all increasing) or they may evolve in different directions (e.g. three increasing and one decreasing). Here, we examine three examples 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.

Example 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 Figure 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.

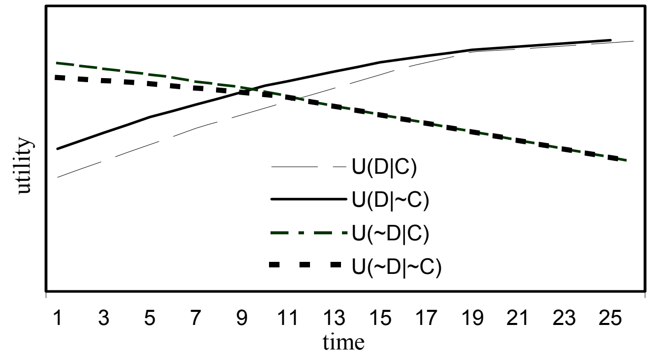

Figure 2. Monotonically increasing and decreasing utilities

Example 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 $U_t(D|C_0)$ are decreasing. This situation can be depicted as in Figure 3. In this case, the decisions are converging approximately at time slice 16 while utilities converge after time slice 22. Hence, this case also supports fast decision convergence.

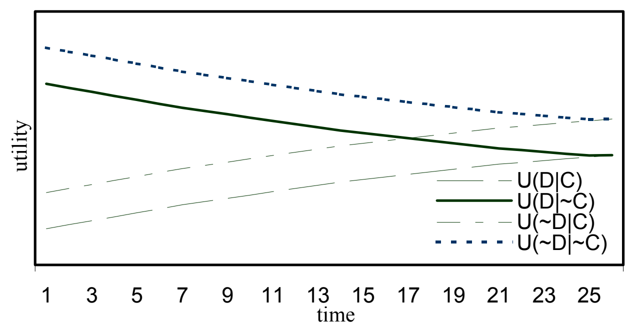

Figure 3. Decision converges faster than utilities

Example 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 takes the same time to converge. An example of this is presented in Figure 4. Here, $U_t(\neg D|C_0)$ and $U_t(\neg D|\neg C_0)$ converge very quickly. On the other hand, $U_t(D|C_0)$ and $U_t(D|\neg C_0)$ take a long time to converge. By looking at the graph we can see that in this case utility and decision both are converging almost at the same time (approximately at time slice 18).

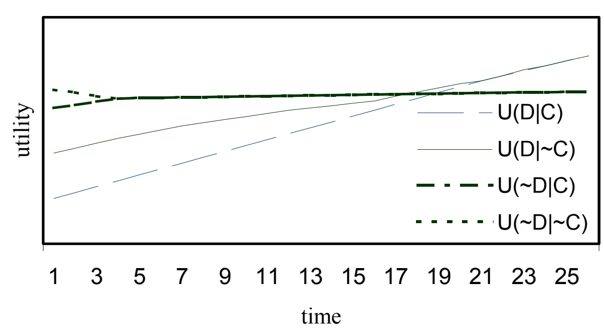

Figure 4. Utilities and decision converge at the same time

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

V. Empirical Generalization

The analytical bound on information relevance in dynamic decision networks derived in Section III describes the behavior of a single variable system. Applying the analytical approach used in Section III to generalized networks results in complex recurrence relations that do not lend themselves to analytical solutions. To generalize the result in Section III 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 [3]. Mathematicians have developed bounds on the convergence of random walks, diffusion models, card shuffles …, etc.

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 presented in Section III

$$T \le \frac{-\log\delta + \log\left|U(D\,|\,C_t) - U(D\,|\,\neg C_t)\right|}{-\log\left|1 - p_1 - p_2\right|}.$$

This bound depend on the chosen accuracy level δ, the values of the rewards $U_t(D|C_t)$ and $U_t(D|\neg C_t)$ 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 for the empirical time bound:

$$T \le \frac{-\log\delta + \log\left|Max(\Delta U)\right|}{-\log\left|Max(\Delta P)\right|}.$$

The notation $Max(\Delta U)$ represents the maximum difference in the utility table and $Max(\Delta P)$ represents the maximum difference in any conditional probability table associated with a temporal edge. To assess the quality of this measure, we have conducted extensive experimental evaluation and the initial results are promising.

VI. Experimental Evaluation

We have conducted several sets of experiments using five dynamic decision network models, namely Generic, Weather Forecast, Car Sales, POMDP and Time Critical. The generic network shown in Figure 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 Figure 5 represents two slices of this network that is slightly more complex than the generic network as the decision node (umbrella) depends on two nodes (forecast and weather). 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 gains. 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 6 shows a time slice of this network. The POMPD network represents a partially observable Markov process with an observable chance node and an unobservable state that determines the reward. Figure 7 depicts a time slice of this network. The time critical network [6], shown in Figure 8 represents a medical intervention scenario.

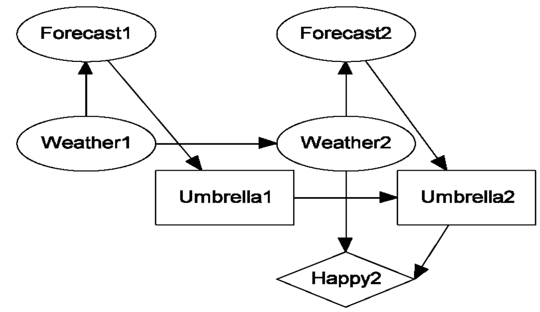

Figure **5.** The Weather Network

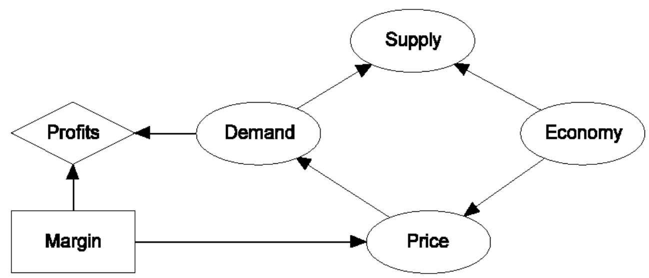

Figure 6. The car sales network

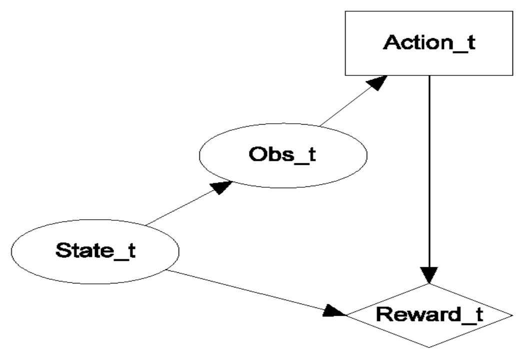

Figure 7. POMDP network

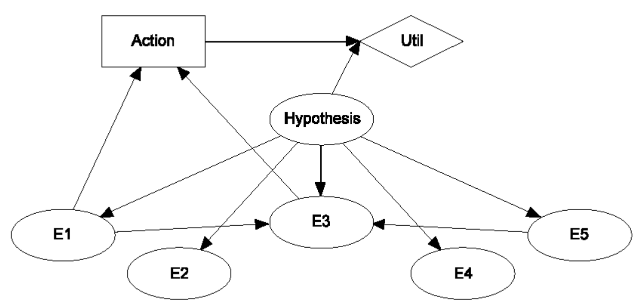

Figure 8. Time-critical Network

A. Utility Convergence Results

In each set of experiments, we use several thousands of randomly generated dynamic decision networks as indicated in Table 1 and Table 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 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 C time slices, where C is a variable that depends on the nature of the network model under consideration. The more complex the network, the less number of time slices are allowed. DDNs that take more than C to converge have been eliminated from the evaluation after verifying that the predicted convergence time exceeds C. This turned out to be true in all test cases.

For the first set of experiments, we use the Generic network with binary valued nodes to show that our proof in the Section III, and our software implementation are indeed correct. The results were 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 where the formula was under-predicted. Moreover, the formula was under-predicted rarely for the Weather Forecast network.

Table 1. Performance of Proposed Predictor

Network	Number of nodes per slice	Delta	Number of tests	Under-Estimated	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%

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 underpredictions were off by 7 or less time slices, a 'slush factor' can be added to the result of the formula to guaranty zero underestimation.

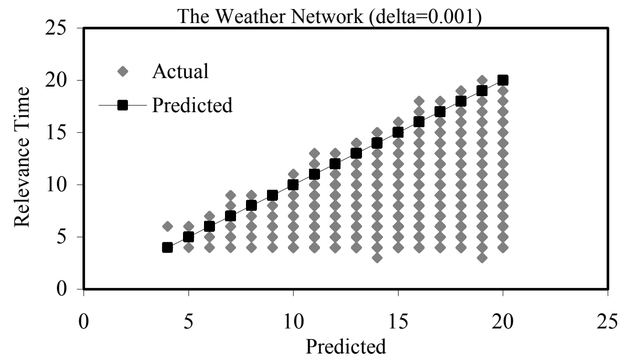

Figure 9. Convergence time of the weather forecast network

Table 2. Relevance Time for Multivalued Networks

Network	Number of States	Delta	Iteration	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%

B. 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.

	Decision	Utility
Max.	12	17
Min.	1	2
Avg.	1.68	4.67
Median	2	4

Table 3. Decision versus Utility Convergence

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.

VII. CONCLUSION

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 computational performance without compromising the quality of the decision.

ACKNOWLEDGMENTS

This work is supported in part by the Natural Sciences and Engineering Research Council (NSERC) Research Grants program and the University of Windsor.

References

[1] Cooper, G. F. (1990). The computational complexity of probabilistic reasoning using Bayesian belief networks, Artificial Intelligence 42(2-3): 393--405.

[2] Dagum, P. , Galper, A., and Horvitz, E. (1992). Dynamic network models for forecasting. In Proceedings of the Eighth Workshop on Uncertainty in Artificial Inteligence, pages 41-48.

[3] Dean, T. and Kanazawa, K.(1989). J. A model for reasoning about persistence and causation. Comp. Int., 5(3).

[4] Dean, T. and Welman, M (1991) Planning and Control, San Mateo, CA: Morgan Kaufmann Publishers.

[5] Diaconis, P. and Shahshahani, M. (1987) Time to reach stationarity in the Bernoulli Laplace diffusion model, SIAM Jour. Math. Anal., 18, 208-218.

[6] Horvitz, E. and Seiver, A. (1997) Time-critical action: Representation and application. In Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI-97).

[7] Huang, T., Koller, D., Malik, J., Ogasawara, G., Rao, B., Russell, S., & Weber, J. (1994). Automated symbolic traffic scene analysis using belief network*s*. In Proceedings of the Twelfth National Conference on Artificial Intelligence, pp. 966-972.

[8] Jensen, F. and Liang, J. 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), 178-183, Paris, France, 1994.

[9] Magni, P.(1999) A new approach to optimal dynamic therapy planning. In Transforming Health Care Through Informatics: Proceedings of the 1999 AMIA Annual Symposium

[10] McCarthy, J., and Hayes, P. J. Some philosophical problems from the standpoint of artificial intelligence. In Machine Intelligence 4, B. Meltzer and D. Michie, Eds. Edinburgh University Press, Edinburgh, Scotland, 1969, pp. 463--502.

[11] Murray, R.C. and VanLehn, K. (2000) DT Tutor: A Decision-Theoretic, Dynamic Approach for Optimal Selection of Tutorial Actions. In G. Gauthier, C. Frasson, & K. VanLehn (Ed.), *Intelligent Tutoring Systems, 5th International Conference, ITS 2000*, pp.153-162. New York: Springer.

[12] Pearl, J. (1988) Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers.

[13] Tawfik, A. Y. and Barrie, T. (2000) The Degeneration of Relevance in Uncertain Temporal Domains: An Empirical Study. *LNAI*, 1822, pages 421-431, Springer-Verlag Berlin Heidelberg.

[14] Tawfik, A. Y. and Neufeld, E. M (2000). Temporal Reasoning and Bayesian Networks. *Computational Intelligence*, 16(3), pages 349-377, Blackwell Publishers, Malden, MA.

[15] Tawfik, A. Y. (2002) Towards temporal reasoning using qualitative probabilities, In Proceedings of FLAIRS-02, the Fifteenth Florida International Artificial Intelligence Research Symposium, Pensacola, FL.