Active Information Fusion For Decision Making Under Uncertainty

Yongmian Zhang
Department of CS
Univ. of Nevada, Reno
Reno, NV 89557, USA
zhang_y@cs.unr.edu

Qiang Ji
Department of ECSE
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
qji@ecse.rpi.edu

Carl G. Looney
Department of CS
Univ. of Nevada, Reno
Reno, NV 89557, USA
looney@cs.unr.edu

Abstract – *Many information fusion applications especially in military domains are often characterized as a high degree of complexity due to three challenges: 1) data are often acquired from sensors of different modalities and with different degrees of uncertainty; 2) decision must be made quickly; and 3) the world situation as well as sensory observations evolve over time. In this paper, we propose a dynamic active information fusion framework that can simultaneously address the three challenges.*

The proposed framework is based on Dynamic Bayesian Networks (DBNs) with an embedded active sensor controller. The DBNs provide a coherent and unified hierarchical probabilistic framework to represent, integrate and infer corrupted dynamic sensory information of different modalities. The sensor controller allows it to actively select and invoke a subset of sensors to produce the sensory information that is most relevant to the current task with reasonable time and limited resources. The proposed framework can therefore provide dynamic, purposive and sufficing information fusion particularly well suited to applications where the decision must be made from dynamically available information of diverse and disparate sources.

To verify the proposed framework, we use target recognition problem as a proof-of-concept. The experimental results demonstrate the utility of the proposed framework in efficiently modeling and inferring dynamic events.

Keywords: Active information fusion, Dynamic Bayesian networks, decision making.

1 Introduction

Many military applications such as target tracking and identification are often characterized as a high degree of complexity and require the multiple perspectives of numerous sensors. Information in the central data fusion side generally involves multiple data types such as various sensory signals, circumstantial evidences, geographical information, subjective knowledge, and various constraints. Moreover, the information collected by different platforms are often uncertain due to sensor noise, target occlusion, perspective distortion, illumination changes and image shadowing. As such, there is a high demand for a fusion system that is able to systematically and efficiently combine, analyze and integrate an enormous volume of information with different degrees of uncertainties and in different levels of abstraction.

In addition, for many military domains, the world situation is often dynamic and unfolds over time. The sensory observations also evolve over time to reflect changes in the world. As a consequence the dynamic aspect of the military domain requires a framework of fusion system to be a time-varying dynamic model that not only captures the beliefs of the current events, but predicts the evolution of different scenarios as well. To correctly assess and interpret the world situation, an adaptive system is therefore needed that not only can systematically handle corrupted sensory data of different modalities but, more importantly, can reason over time as well. The inability of current sensory fusion systems to correlate and reason about a vast amount of information over time is an impediment to providing a coherent overview of the unfolding events since it is often the temporal changes that provide critical information about what we try to infer and understand.

It is crucial for a real time decision support system, especially for tactical defense, to make fast decisions with limited resources. Though the evidences may collect from many possibly allocated sensors with varying capabilities, it is important to avoid unnecessary or unproductive sensor actions and computations. For example, in an air-to-air engagement scenario, an F-15 fighter encounters an incoming aircraft at long range.

The multi-sensor fusion system on-board the F-15, as

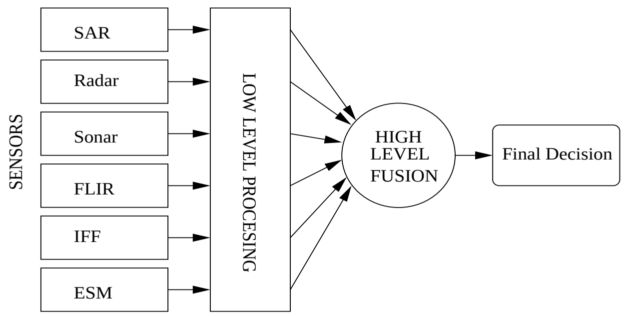

Figure 1: A multi-sensor information fusion system on on-board the flight F15, where SAR - Synthetic Aperture Radar. FLIR - Forward Looking Infra Read; ESM - Electronic Support Measures; and IFF - Identification Friend or Foe.

illustrated in Fig. 1, must quickly assist the pilot to determine the identity of aircraft, i.e., friend or foe. The time and resource constraints mean the sensor fusion system can not activate each possible sensor in order to determine the plane identity. For example, if the attribute measurement oriented sensors, such as Electronic Support Measures (ESM) and Identification Friend or Foe (IFF) in Fig. 1, can yield more reliable information for decision assessment than other sensors in a given time and a given situation, only sensors ESM and IFF then need to be activated at the current time frame. The question is how to determine which information to acquire and which sensor sequence to use to obtain the identified information at a time in order to best answer current query goal. Furthermore, the best information varies with time as the situation changes even if the goal remains unchanged. However, it is our belief that there exists an optimal sensor action strategy that can achieve the goal in two-fold: sufficiency and efficiency. In order to work in a dynamic and time-critical situation, the system must be able to consider different courses of action, and take a subset of most effective sensor strategy which will produce the most informative information to support the identification estimation and to quickly achieve a desired level of confidence to the goal.

In attempting to solve the problems mentioned above, a real time multi-sensory fusion system therefore requires the capability not only to represent the temporal changes in uncertain sensor information, but to dynamically select the most relevant sensory data for a given goal and at a given time as well. There are a number of ways to achieve these goals including the use of the Dynamic Bayesian Networks (DBNs). We select DBNs based on the following considerations. First, DBNs provide a coherent and unified hierarchical probabilistic framework for sensory information representation, integration, and inference [1, 2, 3, 4]. The uncertainties associated with the evidential information can also be incorporated probabilistically in each node as an attribute to characterize the quality of the information. Second, DBNs dynamically evolve and grow to accommodate the new events and to assess the current situation not only based on current information but also to utilize information produced during previous time frames, as alternative scenarios are reinforced or ruled out dynamically. Third, the DBNs can predict the influence of possible future actions on current tasks. This makes it possible to design a sensor action scheme for selectively acquiring information in the light of the data gathered so far.

Using sensor sequential decision making to improve the system performance has recently become an active area of research, especially in robot navigation and computer vision systems [5, 6, 7, 8]. A Bayesian Network (BN) based active fusion for multiple source remote sensing image understanding can be found in [9]. However, all these existing approaches are for static environments. This paper describes some proof-of-concept results of our proposed dynamic active fusion system. We will incrementally incorporate these research results and gradually build a framework that has the capability of simultaneously addressing the following three issues: 1) systematically represent information at different levels of abstraction and with different degrees of uncertainties; 2) account for temporal changes; and 3) perform purposive and sufficing information fusion.

The remainder of this paper is organized as follows. Section 2 gives a brief introduction to DBNs. Section 3 outlines the architecture of the proposed dynamic active information fusion framework. A profit utility for sensor decision making is also derived in this section. Experimental results and analysis are presented in Section 4. We finally provide conclusions and discussion in Section 5.

2 Dynamic Bayesian Networks

To offer the background necessary for us to introduce our approach, in this section we provide a short introduction to DBNs.

Before discussing DBNs, let us first review basic concepts of Static Bayesian Networks (SBNs) [10]. A Bayesian network is a graphical model for representing probabilistic relationships among a set of variables. The network forms a directed acyclic graph (DAG), where nodes represent random variables and directed links between the nodes represent casual relationships. To simplify computation, conditional dependences are systematically built in the Bayesian networks. It is assumed that two nodes at the same level are conditionally independent given their parent node. Also, given the parent node, the child node is independent of the grandparent node. Fig. 2 gives a simple static Bayesian network.

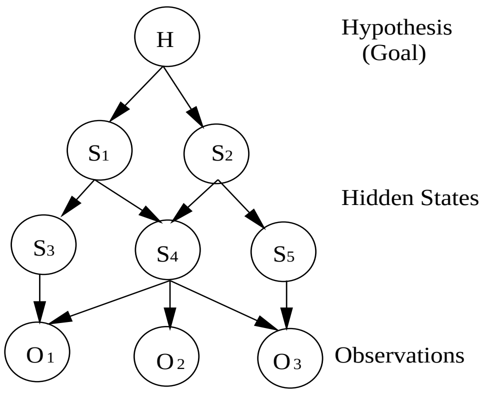

Figure 2: A simple static Bayesian network, where H represents the hypothesis, Ss represent the hidden states/nodes, and Os represent sensory observations.

The SBNs work with evidence and beliefs from a single instant in time. As a result, SBNs are not particularly suited to systems that evolve over time. DBNs are developed to overcome this limitation. DBNs are an extension of SBNs specialized to better model real-world domain with a unified model of time and uncertainty. Significant research has been done in this area [1, 2, 3, 11, 12, 13]. In general, a DBN is made up of interconnected time slices of SBNs, and the relationships between two neighboring time slices are modeled by a Hidden Markov model, i.e., random variables at time t are affected by observable variables at time t, as well as by the random variables at time $t-1$ only. Fig. 3 illustrates such behaviors. The evidence and inferred beliefs of previous time slices are used to estimate and predict beliefs in the current and future events through the causal links, as well as temporal links.

Inference is performed by keeping in memory two slices at any one time, representing the previous discrete time and current time respectively. The slice at the previous time provides diagnostic support for current slice and it is used in conjunction with current sensory data to infer the current hypothesis. The two slices are such programmed that they rotate as old slices are dropped and new slices are used as time progresses. For convenience in belief propagation, at any time instant, the two time slices are treated as an extended SBN and the existing belief propagation for SBN can therefore be applied. There has been substantial research effort concentrated on reducing the computational complexity of performing inference on

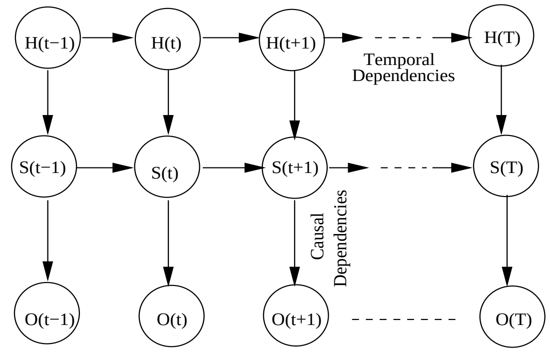

Figure 3: A generic structure for DBN, where $H(t)$ represents hypothesis to the world situation at time t; $S(t)$ represent hidden states at time t; $O(t)$ represents sensory observation at time t; and T is the time boundary.

DBNs [14, 15, 16]. Statistically, the DBN inference can be summarized as follows. Given the DBN topology as shown in Fig. 3, we assume that, the hidden state variables are $S = \{s_0, ..., s_{T-1}\}$, and observable variables are $O = \{o_0, ..., o_{T-1}\}$. Furthermore, let H be the hypothesis variables and $H = \{h_0, ..., h_{T-1}\}$, where T is the time boundary. The inference of hidden variables S by observation O can be expressed as

$$P(S,O) = \prod_{t=1}^{T-1} P(s_t|s_{t-1}) \prod_{t=0}^{T-1} P(o_t|s_t) P(s_0) \qquad (1)$$

and hypothesis H can then be inferred via hidden-state variables S is given by

$$P(H,S) = \prod_{t=1}^{T-1} P(h_t|h_{t-1}) \prod_{t=0}^{T-1} P(s_t|h_t) P(h_0). \qquad (2)$$

The inference is called a bottom-up inference by propagating the states of O to hypothesis H through hidden state variables S. As opposed to the bottom-up inference, by propagating the states of H down to their children nodes, rather than states of observation O, we treat it as a top-down inference. This bi-directional inference mechanism plays a key role in our proposed active fusion framework described in next section.

3 Active Information Fusion

Many real time applications are often constrained by limited time and resources, and decisions must be made fast. Thus, a fusion system that is active, purposive, and sufficing is more suitable for such applications. An active information system is to selectively choose those information sources that are most informative to the problem while minimizing the associated costs in terms of computational complexity, time, and required resources in acquiring the information. Overall efficiency can be achieved by aggregating only a subset of the most relevant sensory data to address current problem. Fig. 4 shows a general view of active information fusion system. The problem that we are interested in is to control (select actions and make decision in) a fusion system that has a repertoire of actions such that the system operates in a purposive manner.

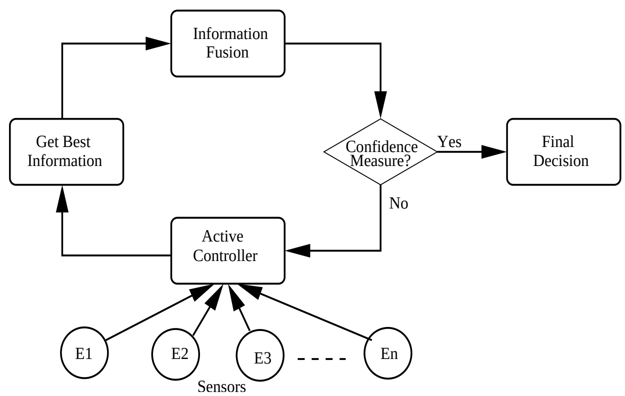

Figure 4: A global overview of active information fusion modules.

Solving the control of selective fusion problem involves several difficult issues: a query/task must be represented in the system and the system must use the representation to operate in a goal-oriented manner. The system must decide what information it needs to gather next to maximize its solution to the goal. It also needs to evaluate the cost/benefit of a sensory action for a given task. Finally, it needs to quantitatively evaluate and characterize the quality of the information.

DBNs are appropriate for performing selective and sufficing sensory fusion since DBNs provide both top-down and bottom-up inference mechanisms. The top-down inference can be used to predict the utility of a particular sensory action with respect to a goal at the top, the bottom-up inference allows the integration of the sensory information from a sensory action and to update beliefs for each node in the network. Fig. 5 provides a pictorial summary of a DBN based framework for active information fusion system.

To obtain the desired information, the information module needs to determine the next optimal information and sensory action to perform. In the language of decision theory, the selection of an information source or the activation of a process to compute new information are simply regarded as a set of actions available to the decision maker. Choosing an action will have a consequence (e.g., altering the confidence in a hypothesis variable and costing resources). If we can devise

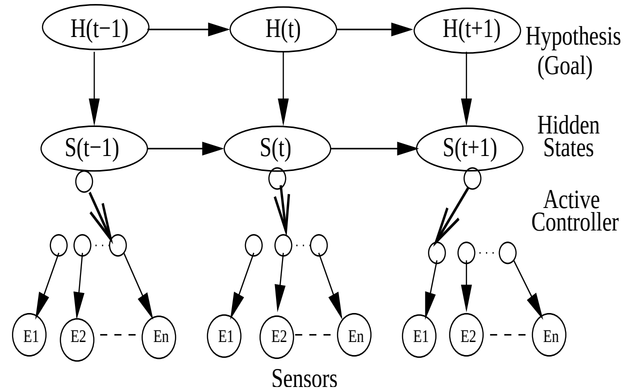

Figure 5: Dynamic Bayesian Networks based active information fusion framework. The system consists of a Goal, Hidden States, an Active Sensor Controller, and numerous of Information Sensors.

a benefit/cost measure to each possible consequence, this can be used by the system to decide what action to perform, and what sensor to activate. The consequences of an action are an uncertain quantity. So the best thing the decision maker can do is to choose the action that could maximize its expected utility. The expected utility depends on the current available sensory data, the internal knowledge, and the current goal. In target recognition, as an example, different sensors may be activated based on their usefulness on the information content for recognizing a particular type of target and on the information already collected so far.

Now, assume that we have the sensory information sources $\mathbf{E} = \{E_1, ..., E_m\}$, which is a set of measurements taken from sensors labeled from 1 to m respectively, and let H be the hypothesis to confirm, as shown in Fig. 5. The most informative evidence is the one that decreases the uncertainty of the hypothesis H the most. This means that we will give high preference to the sensor that can lead the probabilities of hypothesis close to near one and zero. Let V be the uncertainty-reducing potential to the hypothesis, and given a piece of evidence e from a sensor E_i, which has n_e possible value, the most expected value of V for hypothesis h with n_j states, can be defined as

$$V_i = \max_{k=0}^{n_e} \sum_{j=0}^{n_h} [P(h_j|e_k)]^2 - \min_{k=0}^{n_e} \sum_{j=0}^{n_h} [P(h_j|e_k)]^2, \quad (3)$$

where i is a sensor tag which identifies the sensor that supplies the evidence information and $i \in \{1, 2, ..., m\}$; $P(h_j|e_k)$ can be obtained by propagating the possible outcome of an information source, i.e

$$P(h_j|e_k) = \frac{P(h_j, e_k)}{P(e_k)} \quad (4)$$

Access of information requires cost such as the cost of information retrieval, time delay and extra computation time, etc., even though it is certainly available. If we consider the sensor costs in a general term, the expected profit utility can be expressed as

$$u_i = \alpha V_i + (1-\alpha)(1 - \frac{C_i}{\sum_{i=1}^{n} C_i}), \qquad (5)$$

where C_i is the cost to acquire the information by sensor with tag i, and it usually needs domain knowledge; α denotes the balance coefficient between the belief of goal and the cost of information acquisition, and $0 \leq \alpha \leq 1$. Equation 5 compromises between contribution to the belief of goal and the cost associated with acquiring these information sources to achieve the desired level of confidence to the goal. From equation 5 we can see that the utility value of an evidence increases with the belief accorded to the goal by the evidence and decreases as the cost to acquire the evidence increases. We can choose an optimal sensor action A^* using following decision rule

$$A^* = \arg\max_{A} \sum_{j} u(E \subset \mathbf{E}, h_j) P(h_j | E \subset \mathbf{E}). \qquad (6)$$

$E \subset \mathbf{E}$ is a subset of observation for a set of sensory information sources $\mathbf{E}$ at a particular point of time. If we incorporate time dimension to the modeling event, the probability distribution of the goal we want to achieve can be generally described as

$$P(H, A^*) = \prod_{t=1}^{T-1} P(S_t|S_{t-1}) \prod_{t=0}^{T-1} P(H_t|E_t) P(H_0), \qquad (7)$$

where T is time boundary; the hypothesis $H = \{H_0, ..., H_t, ..., H_T\}$ and the subset of sensed information $E = \{E_0 \subset \mathbf{E}, ..., E_t \subset \mathbf{E}, ..., E_T \subset \mathbf{E}\}$, on time sequence of T. S represents a set of any hidden states including hypothesis that has temporal links between corresponding nodes in two neighboring time frames, and the set of hidden states $S = \{S_0, ..., S_t, ..., S_T\}$ on time sequence of T. Normally the best action varies with time. The sensor action strategy therefore must be recalculated at each time slice. The algorithm is summarized as follows:

1. Compute uncertainty-reducing potential $\{V_1^t, .., V_m^t\}$ by Eqn. 3
2. Calculate utility value $\{u_1^t, ..., u_m^t\}$ using Eqn. 5
3. Select the most informative subset of information source $E \subset \mathbf{E}$ based on Eqn. 6
4. Instantiate the subset of information sources E
5. Run DBN inference algorithm to update belief
6. If the goal hypothesis belief $P(H, A^*) \geq$ confidence threshold, then terminate; otherwise
7. Add a new time slice $t = t + 1$, and go to step 1

4 Experiments

This section describes simulations performed to verify the proposed active fusion framework. We experimentally analyze the performance of active and passive fusion in order to illustrate how the proposed active fusion system basically works. The test case we present is a target recognition problem, in which an air-to-air engagement scenario is simulated.

4.1 Problem Definition

In an air-to-air engagement scenario, an F-16 fighter encounters an aircraft, the multi-sensor fusion system is activated and a goal is to determine the identity of the aircraft, i.e., Friend or Foe. We assume that the aircraft is equipped with four on-board sensors for identifying the type of interrogated aircraft as shown in Fig. 1 of Section 1. The sensors may include imaging sensor (e.g., FLIR), acoustic sensor, and radar sensor, etc., to identify different attributes of the aircraft such as shape, color, motion, speed, and sound. For the task of identifying different types of aircrafts the lowest level of abstraction represents the observable sensory data. The intermediate level contains the identification of different types of aircrafts e.g., F-15, B-2, MiG-29. Intermediate levels of abstraction, i.e., different partitions of the aircraft types, are: nature which has subdivisions of friend, foe, and neutral; and class which is often divided into fighter, bomber, and airliner. Each of these alternate levels of abstraction may be the focus of a sensor available to the multi-sensor data fusion system. The target node is the aircraft type. There is also contextual information such as circumstantial evidences and geographical information, which may contribute to the presence of certain type of aircraft. Fig. 6 shows a BN modeling the problem of aircraft type identification. The network shown in Fig. 6 merely represents the static component of the DBN modeling. The interconnections among different static components by HMM forms the complete topology of DBN for aircraft recognition as shown in Fig. 7.

The time constraints do not allow the fusion system to activate all possible sensors in a passive manner. We therefore need to determine the best sensory action scheme to accelerate identification estimation. For the aircraft type determination task, we need the attributes that best characterize the aircraft type. The attributes may include shape, speed, color, special markings, and sound. Based on the expected utility

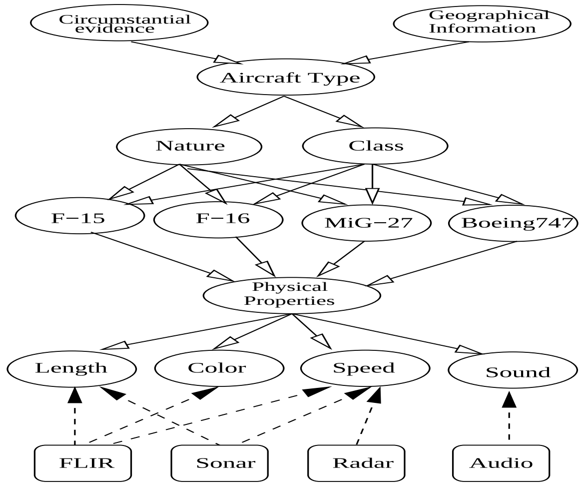

Figure 6: The topology of BN for target recognition. The network merely represents the static component of the DBN modeling. Note: FLIR, Sonar, Radar, Audio are not the part of network nodes, but information sources.

computation, shape and speed are the most important attributes. To obtain this information, we may need to activate the imaging sensors to determine the shape and motion parameters of the aircraft. For speed, we can also use the sonar sensor. The choice of which sensor to activate depends on the expected utility of each sensor. After activating the sensor, the information we obtain on the shape and speed of the aircraft can help determine the aircraft type. If we still are not confident at the aircraft type, we need to decide what will be the next sensor action at the next time slice. This repeats until we identify the aircraft type with sufficient confidence. The application domain can be formulated as a sequential stochastic Markov decision process.

4.2 Experimental Results

The attributes of an interrogated aircraft are assumed to be extracted by low level imaging and non-imaging data processing modules. The numeric attributes are then fuzzified by mapping them into fuzzy classes. The input sensory data into the fusion framework is finally to be a fuzzy descriptor, e.g., using "very long", "long", "average", and "short" to describe the length of aircraft. The initial probability distributions as well as transition distributions between time slices are assigned subjectively based on experiences in this test. Due to space limitation, the prior and conditional probabilities are not listed in this paper. The transitional probabilities between corresponding nodes in two neighboring slices are also subjectively determined and are listed in Table 1. As a part of profit utility, the acquisition cost must be determined using domain knowledge of specific sensors. However, we subjectively assign them in Table 2 for this particular test. Fig. 7 presents the temporal links of intermediate nodes between time $t-1$ and t. In this particular case we assume neither circumstantial evidences nor geographical information are relevant for real-time fusion on-board the aircraft. The preference weight α in Equation 5 takes 0.6 when the information acquisition cost is considered.

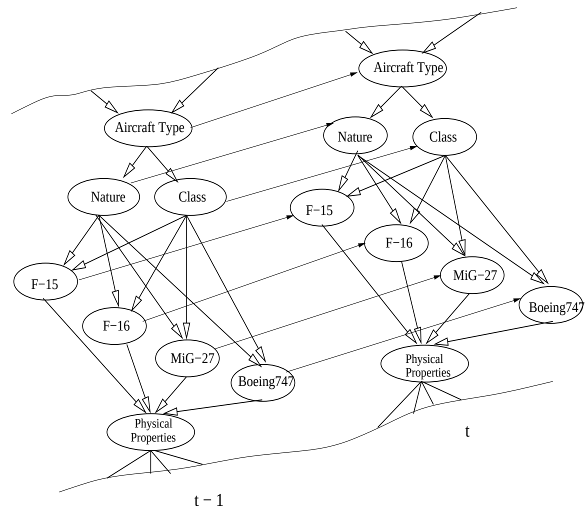

Figure 7: The dynamic model for target recognition problem. This figure only shows nodes have temporal links between neighboring time slices. Note that Physical Properties is not a part of dynamic linked nodes.

Table 1: Transition probabilities among time slices

Var.	AirType	Nature	Class	F15	F16	MG27	747
Pr.	1.0	0.9	0.8	1.0	1.0	1.0	1.0

Table 2: Sensory information acquisition cost

Sensor	FLIR	Sonar	Radar	Audio
Acquisition cost	5	6	4	7

The results given in Table 3 and Table 4 demonstrate how the sensor active controller takes action during information fusion. Table 3 gives the result when the acquisition cost is ignored; while Table 4 provides the result with consideration of information acquisition cost. Table 5 and Fig. 9 show the results of fusing two sensors at each time slice. The activation sequence for the passive fusion is randomly generated and the value plotted in Fig. 8 and fig. 9 is averaged over 5 runs.

We observed that, in Table 3, besides the sensor $FLIR$ and $Sonar$ instantiated at time slice 3 and 4, $Radar$ dominates at the rest of time slices. In other

Table 3: Sensor decision making without considering information acquisition cost

Time Slice	Sensor Selected	Profit Utility Value			
		FLIR	Sonar	Radar	Audio
1	Radar	0.0601	0.0368	0.2779	0.2134
2	Radar	0.2430	0.2438	0.3034	0.1906
3	FLIR	0.3489	0.3326	0.2796	0.1601
4	Sonar	0.2979	0.3062	0.3007	0.1694
5	Radar	0.2560	0.2644	0.3058	0.1743
6	Radar	0.1982	0.2049	0.2673	0.1524
7	Radar	0.1433	0.1466	0.2386	0.1375
8	Radar	0.0968	0.0982	0.2195	0.1283

Table 4: Sensor decision making with consideration of information acquisition cost

Time Slice	Sensor Selected	Profit Utility Value			
		FLIR	Sonar	Radar	Audio
1	Audio	0.1270	0.1312	0.2395	0.2553
2	Radar	0.1839	0.2191	0.2439	0.2385
3	Sonar	0.2867	0.2979	0.2349	0.2205
4	Sonar	0.2607	0.2844	0.2393	0.2207
5	Sonar	0.2412	0.2633	0.2423	0.2235
6	Radar	0.2153	0.2379	0.2438	0.2304
7	Radar	0.1845	0.2045	0.2252	0.2205
8	Audio	0.1557	0.1746	0.2125	0.2142

words, system attempts to repeatedly fuse speed and length of the target platform just from a single *Radar*. Nevertheless, as shown in $Curve1$ of Fig. 8, this can lead to more quickly achieve the goal than the passive fusion does. Practically, we need to avoid information over-redundancy [1] . This is apparently the case if a Radar reports the same value of target speed to a fusion system for hours at time slice of only a few seconds. However, merging appropriate redundancy of sensory information may be beneficial for reducing imprecision and increasing reliability.

Table 4 shows that active sensors alternate more frequently as the acquisition cost is accounted for, and as such, the benefits of active fusion of multiple sensors will not lose. Furthermore, we can see from $Curve2$ of Fig. 8 that there is no significant performance difference between that with and without the consideration of the acquisition cost. This also indicates that, if we are able to devise an appropriate benefit/cost model as well as active sensor controller, over-redundancy of sensory information can be overcome. The similar performance can be seen from Table 5 and Fig. 9 as two sensors are simultaneously activated at each time slice. The results are particularly apparent from Fig. 8 and Fig. 9 that the active fusion outperforms the passive fusion in term of time spent to achieve the goal.

5 Conclusions

Many information fusion applications especially in military domains are often characterized as a high de-

[1] *over-redundancy* refers to repeated use of the same sensor in consecutive time frames.

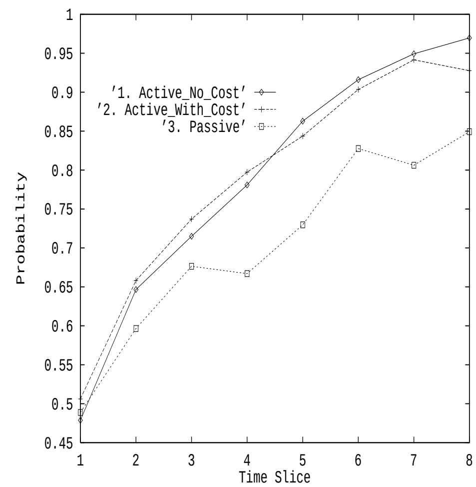

Figure 8: The simulation results of target recognition problem. Only one sensor is activated at each time slice. $Curve1$ - active fusion with acquisition cost ignored; $Curve2$ - active fusion with the information acquisition cost considered; $Curve3$ - passive information fusion.

Table 5: Sensor decision making with acquisition cost (two sensors are active at each time slice)

Time Slice	Sensor Selected	Profit Utility Value			
		FLIR	Sonar	Radar	Audio
1	Audio-Radar	0.1270	0.1312	0.2395	0.2553
2	Sonar-Radar	0.2367	0.2554	0.2547	0.2416
3	Sonar-FLIR	0.3003	0.3087	0.2405	0.2234
4	Sonar-FLIR	0.2697	0.2928	0.2531	0.2289
5	Sonar-Radar	0.2453	0.2690	0.2657	0.2379
6	Radar-Sonar	0.2118	0.2344	0.2420	0.2242
7	Radar-Audio	0.1883	0.2101	0.2419	0.2297
8	Radar-Audio	0.1725	0.1959	0.2786	0.2686

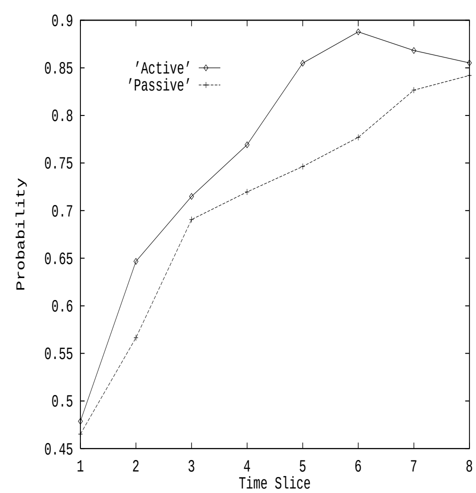

Figure 9: The simulation results of target recognition problem. There are two sensors are active at each time slice.

gree of complexity due to three challenges: 1) sensory information obtained from multiple perspectives of sensors is often corrupted; 2) decisions must be made quickly; and 3) the world situation as well as sensory observations evolve over time.

In this paper, we described a probabilistic framework based on DBNs to simultaneously address the three challenges. The proposed system is based on DBNs embedded with an active sensor controller. The active sensor controller allows it to actively select and invoke a subset of sensors to produce the sensory information that is most relevant to current task with reasonable time and limited resources. The proposed framework can provide dynamic, purposive and sufficing information fusion particularly well suited to time-varying and time critical fusion applications.

In summary, we believe that the major contributions of this work is to present a general framework that can be used in a variety areas where decisions must be made quickly and economically from dynamically available information of diverse and disparate sources. Future work will involve incremental improvements of the framework and application of it to more complex problems.

Acknowledgments

This material is based upon work supported by a grant from the US Army Research Office under Grant No. DAAD19-01-1-0402.

References

[1] G. M. Provan and J. R. Clarke, "Dynamic network construction and updating techniques for the diagnosis of acute abdominal pain," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 15, no. 3, pp. 299–307, 1993.

[2] A. Nicholson and J. Brady, "Dynamic belief networks for discrete monitoring," *IEEE Trans. on System, Man and Cybernetics*, vol. 24, no. 11, pp. 1593–1609, 1994.

[3] V. Pavolovic, *Dynamic bayesian networks for information fusion with applicaiton to human-computer interfaces.* PhD thesis, University of Illinios at Urbana-Champaign, 1999.

[4] T. Dean and M. P. Wellman, *Planning and Control.* New York: Morgan Kaufman Publisher, 1991.

[5] R. D. Rimey and C. M. Brown, "Control of selective perception using bayes nets and decision theory," *International Journal of computer Vision, Special Issue on Actice Vision*, vol. 12, no. 2-3, pp. 173–207, 1994.

[6] T. Levitt, T. Binford, G. Ettinger, and P. Gelband, "Probability-based control for computer vision," in *Image Understanding Workshop*, pp. 355–369, 1989.

[7] S. A. Hutchinson and A. C. Kak, "Planning sensing strategies in robot work cell with multi-sensor capabilities," *IEEE Journal of Robotics and Automation*, 1989.

[8] T. Dean, T. Camus, and J. Kirman, "Sequential decision making for active perception," in *DARPA Image Understanding Workshop*, pp. 889–894, 1990.

[9] A. Pinz, M. Prantl, H. Ganster, and H. K. Borotschnig, "Active fusion - a new method applied to remote sensing image interpretation," *Pattern Recognition Letters*, vol. 17, no. -, pp. 1349–1359, 1996.

[10] J. Pearl, *Probabilistic Reasoning in Intelligent Systems.* San Mateo, CA: Morgan Kaufmann Publishers, 1988.

[11] T. Dean and K. Kazanawa, "A model for reasoning about persistance and causation," *Computational Intelligence*, vol. 5, no. -, pp. 142–150, 1989.

[12] P. Dagum, A. Galper, and E. Horvic, "Dynamic network models for forecasting," in *Proc. of the 8th Conference on Uncertainty in Artificial Intelligence*, pp. 41–48, 1992.

[13] J. D. Young and E. Santos, "Introduction to temporal networks," in *The 7th Midwest AI and Cognitive Science Conference*, April 1996.

[14] K. Kanazawa, D. Koller, and S. Russell, "Stochastic simulation algorithms for dynamic probabilistic networks," in *Proc. of 11th annual conference on uncertainty in AI*, pp. 346–351, 1995.

[15] N. Jitnah and A. E. Nicholson, "Arc weights for approximate evaluation of dynamic belief networks," in *Proc of the 12th Australian Joint Conference on Artificial Intelligence, AI'99*, pp. 393–404, 1999.

[16] D. Koller and X. Boyen, "Exploiting the architecture of dynamic systems," in *Proceedings of the 16th National Conference on Artificial Intelligence (AAAI-99)*, pp. 313–320, July 1999.