Efficient Active Fusion for Decision-making via VOI Approximation

Wenhui Liao and Qiang Ji
Department of Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute, Troy, NY 12180-3590, USA
{liaow, jiq}@rpi.edu

Abstract

Active fusion is a process that purposively selects the most informative information from multiple sources as well as combines these information for achieving a reliable result efficiently. This paper presents a general mathematical framework based on Influence Diagrams (IDs) for active fusion and timely decision making. Within this framework, an approximation algorithm is proposed to efficiently compute non-myopic value-of-information (VOI) for multiple sensory actions. Meanwhile a sensor selection algorithm is proposed to choose optimal sensory action sets efficiently. Both the experiments with synthetic data and real data from a real-world application demonstrate that the proposed framework together with the algorithms are well suited to applications where the decision must be made efficiently and timely from dynamically available information of diverse and disparate sources.

Introduction

Active fusion extends the paradigm of information fusion, being not only concerned with the methodology of *combining* information, but also introducing mechanisms in order to *select* the information sources to be combined (Pinz *et al.* 1996). By purposively choosing an optimal subset from multiple information sources and fusing these information, it can save computational time and physical cost, reduce redundancy, and increase the chances of making correct decisions. Due to these benefits, it plays an especially important role for timely and effective decision-making.

To model active fusion for timely decision-making, we explore a decision-theoretic paradigm, namely Influence Diagram (ID), to actively fuse information by viewing it as a decision-making problem. Such a model provides a fully unified hierarchical probabilistic framework for representation, integration, and inference of uncertain sensory information of different modalities at multiple levels of abstraction. In addition, it embeds both the sensors' contributions to decision-making and their operating cost in one framework to choose the optimal sensory action set by exploiting utility theory and probability theory.

Within the framework, one key challenge is to efficiently decide an optimal sensory action set. Two issues impede this target: 1) the computation of sensory action selection criteria could have exponential time complexity; 2) the number of sensory action subsets grows exponentially as the total number of sensory actions.

For the first issue, to rate the usefulness of various sensors, a popular criterion is Value-of-Information (VOI), which is defined as the difference of the maximum expected utilities with and without observing certain information in an ID. It is one of the most useful sensitivity analysis techniques for decision analysis. In the past several years, a few methods have been proposed to compute VOI in IDs (Dittmer & Jensen 1997; Shachter 1999; Zhang, Qi, & Poole 1993; Poh & Horvitz 1996). However, they only focus on computing myopic VOI for each individual sensor, instead of non-myopic VOI for a sensor set. In fact, it is usually too cumbersome to compute non-myopic VOI for any practical use since the number of the observation sequences grows exponentially as the number of available sensors. Heckerman et. al. (Heckerman, Horvitz, & Middleton 1993) may be the only one that proposed a solution to this problem. However their method assumes that all the random nodes and decision nodes are binary, and the sensors are conditionally independent given the hypothesis. Motivated by their method, we propose an approximation algorithm to compute non-myopic VOI efficiently by exploiting the central-limit theory. Especially, it doesn't need the assumptions made by Heckerman's method. The efficiency of the algorithm makes it feasible to various applications where efficiently evaluating a large amount of information sources is necessary.

For the second issue, in practice, most work (Wang *et al.* 2004; Ertin, Fisher, & Potter 2004) selects the best sensor myopically, or select the first m sensors by simply ranking individual sensors. But the selected sensors could lead to bad results since the selection methods cannot provide performance guarantees. Other groups (Kalandros, Pao, & Ho 1999; Fassinut-Mombot & Choquel 2004) develop more complex algorithms for achieving global optimal solutions such as searching-based approach. Such algorithms could suffer from speed and end up at local optimum. We present an improved greedy approach for sensory action set selection. Together with the approximation algorithm for VOI computation, it reduces the time complexity significantly.

Overall, with the proposed framework and the computational algorithms, timely and effective decisions can be

made based on a selective set of sensors with the optimal trade-off between their cost and benefit.

Active Fusion Modeling

An Influence Diagram, as shown in Figure 1, is presented to model active fusion as a decision-making problem. The top node Θ represents the target hypothesis variable. It can also be generalized to indicate a set of hypothesis variables. Each bottom node O_i indicates the possible observations from each sensor. Each node I_i associated with O_i represents the sensory information measured by O_i, and the link between them reflects the reliability of sensor measurement. The big ellipse indicates all the chance nodes between the Θ node and I nodes. These nodes are collectively called hidden nodes. They model the probabilistic relationships between the Θ node and I nodes at different abstraction levels.

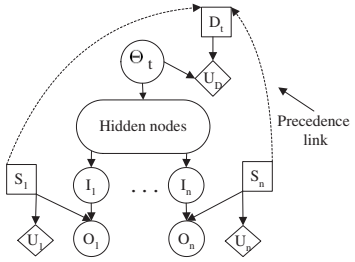

Figure 1: A dynamic ID model for active information fusion. Ellipses denote chance nodes, rectangles denote decision nodes, and diamonds denote utility nodes.

There are two types of decision nodes. One is the decision node D indicating the possible actions associated with the hypothesis node Θ. Another type is the sensory action node S_i. It controls whether to activate a sensor for collecting observation O_i or not. The arcs connecting each sensory node S_i to the decision node D indicate the time precedence order. Corresponding to the decision nodes, there are two types of utility nodes. The utility node connected with both node Θ and D indicates the benefit (penalty) of taking appropriate (inappropriate) actions with respect to a particular hypothesis state. Each utility node U_i connected with a sensory node defines the cost of obtaining data from the sensor.

Such a model provides a coherent and fully unified hierarchical probabilistic framework for realizing three main functions: deciding a sensory action set that achieves optimal trade-off between the cost and benefit of sensors, applying a fusion model for efficient combining the information from the sensor set, and making decisions based on the fusion results. We focus on the first function in this paper.

Non-myopic VOI Computation

In this section, we present the approximation algorithm for non-myopic VOI computation based on the ID framework.

Value-of-Information

As one of the most useful sensitivity analysis techniques of decision analysis, VOI is used to rate the usefulness of various information sources and to decide whether pieces of evidence are worth acquiring before actually using the information sources (Pearl 1988). Let $O = \{O_1, ..., O_n\}$ be a subset of observation variables, the value-of-information of O, $VOI(O)$, can be defined as:

$$VOI(O) = EU(O) - Co - EU(\bar{O}) \quad (1)$$

$$EU(O) = \sum_{O} p(o) \max_{D} \sum_{\Theta} p(\theta|o)u(\theta, d) \quad (2)$$

$$EU(\bar{O}) = \max_{D} \sum_{\Theta} p(\theta)u(\theta, d) \quad (3)$$

where $u()$ denotes the utility function, $EU(O)$ denotes the expected utility to the decision maker if O were observed, $EU(\bar{O})$ denotes the expected utility to the decision maker without observing O, and $Co = \sum_i u(S_i)$, which denotes the total cost to collect observations in O.

As shown in Equation 1, to compute $VOI(O)$, it is necessary to compute $EU(O)$ and $EU(\bar{O})$ respectively. Obviously, $EU(\bar{O})$ is easy to compute, whereas directly computing $EU(O)$ could be cumbersome. If the decision maker has the option to observe a subset of observations $\{O_1, ..., O_n\}$ and each O_i has m possible values, then there are m^n instantiations of the observations in this set. Thus, to compute $EU(O)$, there are m^n inferences to be performed. It makes it infeasible to compute the VOI when n is not small.

It is noticed that each instantiation of O corresponds to a specific optimal action for the decision node D. Therefore we can divide all the instantiations of O into several subsets, where the optimal action is the same for those instantiations in the same subset. Specifically, if D has q decision rules, $\{d_1, ..., d_q\}$, all the instantiations of O can be divided into q subsets, o_{d_1}, ... ,o_{d_q}, where $d_k = \arg\max_D \sum_\Theta p(\theta|o)u(\theta, d_k)$ for $o \in o_d$. Thus, from Equation 1, $EU(O)$ can be further derived as:

$$\begin{aligned} EU(O) &= \sum_{O} \max_{D} \sum_{\Theta} p(\theta)p(o|\theta)u(\theta, d) \\ &= \sum_{\Theta} p(\theta) \sum_{D} \sum_{o \in o_{d_k}} p(o|\theta)u(\theta, d_k) \end{aligned} \quad (4)$$

Therefore, the key is to compute $\sum_{o \in o_{d_k}} p(o|\theta)$ efficiently.

Decision Boundaries

To compute $\sum_{o \in o_{d_k}} p(o|\theta)$, first we need to know how to divide the instantiations of O into the q subsets, o_{d_1}, ..., o_{d_q}. We fist focus on the case that Θ has only two states, θ, $\neg\theta$, and then extend it to the general case later.

Based on the definition, the expected utility of taking action d_k is $Eu(d_k) = p(\theta) * u_{1k} + p(\neg\theta) * u_{2k}$, where $u_{1k} = u(\theta, d_k)$, and $u_{2k} = u(\neg\theta, d_k)$. If d_k is the optimal action, $Eu(d_k)$ must be larger than or equal to the expected utility of any other action:

$$Eu(d_k) \geq Eu(d_j), \text{for} \forall j, j \neq k \quad (5)$$

$$\Rightarrow \quad p(\theta) * u_{1k} + p(\neg\theta) * u_{2k} \geq p(\theta) * u_{1j} + p(\neg\theta) * u_{2j} \quad (6)$$

$$\Rightarrow \quad p(\theta) \geq r_{jk}, \text{if } j > k; \text{otherwise}, p(\theta) < r_{jk}$$

$$\text{where } r_{jk} = \frac{u_{2k} - u_{2j}}{u_{1j} - u_{1k} + u_{2k} - u_{2j}} \quad (7)$$

We assume $u_{1k} > u_{1j}$, $u_{2k} < u_{2j}$ if $k < j$, and $u_{1j} - u_{1k} + u_{2k} - u_{2j} \neq 0$ without loss of generality. Thus, based

on the above equation, the optimal action is d_k iff $p(\theta) \geq \max_{k<i\leq q} r_{ik}$ and $p(\theta) < \min_{1\leq j<k} r_{jk}$. Let the lower threshold for d_k be $p^*_{kl} = \max_{k<i\leq q} r_{ik}$, and the upper threshold be $p^*_{ku} = \min_{1\leq j<k} r_{jk}$, then, if $p^*_{kl} \leq p(\theta) < p^*_{ku}$, d_k is the optimal decision. If $k = 1$, $p^*_{ku} = 1$.

Therefore, $\mathrm{p}(p^*_{kl} \leq p(\theta|o) < p^*_{ku})$ actually indicates the probability that d_k is the optimal decision given certain observation o. On the other hand, $\sum_{o\in o_{d_k}} p(o|\theta)$ is the sum of the probability of each $o \in o_{d_k}$ being observed given θ and each o makes d_k be the optimal decision. Thus, $\sum_{o\in o_{d_k}} p(o|\theta)$ also indicates the expected probability of d_k being the optimal action given some observations.

Therefore,

$$\sum_{o\in o_{d_k}} p(o|\theta) = p(p^*_{kl} \leq p(\theta|o) < p^*_{ku}) \tag{8}$$

Thus, the problem of computing $\sum_{o\in o_{d_k}} p(o|\theta)$ transfers to the problem of computing $p(p^*_{kl} \leq p(\theta|o) < p^*_{ku})$.

Approximation with Central Limit Theorem

To compute $p(p^*_{kl} \leq p(\theta|o) < p^*_{ku})$, one way is to treat $p(\theta|o)$ as a random variable. If such a random variable satisfies certain standard probability distributions, it will be easy to compute $p(p^*_{kl} \leq p(\theta|o) < p^*_{ku})$. However, it is hard to get such a distribution directly. But we notice that $p(p^*_{kl} \leq p(\theta|o) < p^*_{ku}) = p(\frac{p^*_{kl}}{1-p^*_{kl}} \leq \frac{p(\theta|o)}{p(\neg\theta|o)} < \frac{p^*_{ku}}{1-p^*_{ku}})$. In the next, we show that $\frac{p(\theta|o)}{p(\neg\theta|o)}$ can be approximated with a Log normal distribution.

If all the O_i nodes are conditionally independent with each other given Θ, based on the chain rule:

$$\frac{p(\theta|O)}{p(\neg\theta|O)} = \frac{p(O_1|\theta)}{p(O_1|\neg\theta)}...\frac{p(O_n|\theta)}{p(O_n|\neg\theta)}\frac{p(\theta)}{p(\neg\theta)} \tag{9}$$

However, usually some $O_i s$ may not be conditionally independent given Θ. To be able to get the similar format to Equation 9, we first divide O into several groups, where the nodes in one group are conditionally independent with the nodes in other groups. Based on the d-separation concept (Jensen 2001) in Bayesian Networks, O_i and O_j are conditionally independent given Θ if Θ is the only common ancestor that is not instantiated for O_i and O_j. Therefore, with the grouping procedure, O can be automatically divided into several sets, named O^{s1}, O^{s2}, ..., O^{sg}, where g is the number of the groups. Thus, Equation 9 can be modified as:

$$\frac{p(\theta|O)}{p(\neg\theta|O)} = \frac{p(O^{s1}|\theta)}{p(O^{s1}|\neg\theta)}...\frac{p(O^{sg}|\theta)}{p(O^{sg}|\neg\theta)}\frac{p(\theta)}{p(\neg\theta)}$$

$$\Rightarrow ln\frac{p(\theta|O)}{p(\neg\theta|O)} = \sum_{i=1}^{g} ln\frac{p(O^{si}|\theta)}{p(O^{si}|\neg\theta)} + ln\frac{p(\theta)}{p(\neg\theta)}$$

$$\Rightarrow ln\phi = \sum_{i=1}^{g} w_i + c,$$

where $\phi = \frac{p(\theta|O)}{p(\neg\theta|O)}$, $w_i = ln\frac{p(O^{si}|\theta)}{p(O^{si}|\neg\theta)}$, $c = \ln\frac{p(\theta)}{p(\neg\theta)}$ (10)

In the above equation, c can be regarded as a constant reflecting the state of Θ before any new observation is obtained and any new decision is taken.

From the previous analysis, we know d_k is the optimal action with the probability $p(\frac{p^*_{kl}}{1-p^*_{kl}} \leq \phi < \frac{p^*_{ku}}{1-p^*_{ku}})$, which is equivalent as $p(p^*_{kl} \leq p(\theta|o) < p^*_{ku})$. Let $\phi^*_{kl} = \frac{p^*_{kl}}{1-p^*_{kl}}$, and $\phi^*_{ku} = \frac{p^*_{ku}}{1-p^*_{ku}}$, thus d_k is the optimal decision if and only if $\phi^*_{kl} \leq \phi < \phi^*_{ku}$. Thus, the following equations stand:

$$\sum_{o\in o_{d_k}} p(o|\theta) = p(\phi^*_{kl} \leq \phi < \phi^*_{ku}|\theta) \tag{11}$$

$$\sum_{o\in o_{d_k}} p(o|\neg\theta) = p(\phi^*_{kl} \leq \phi < \phi^*_{ku}|\neg\theta) \tag{12}$$

To compute the probabilities in the above equations, the probability distribution of ϕ must be known. Let $W = \sum_{i=1}^{g} w_i$, the sum of w_i. If W satisfies a Gaussian distribution, ϕ will be Log normal distribution based on Equation 10. The central limit theorem states that the sum of independent variables approaches a normal distribution when the number of variables becomes large. Also, the expectation and variance of the sum is the sum of the expectation and variance of each individual random variable. Thus, regarding each w_i as an independent variable, W then follows a Gaussian distribution. In the following, we first compute the mean and variance of each w_i so that to get the distribution of W, then obtain the distrubution for ϕ so that the probabilities in Equation 11 can be easily achieved.

Assume O^{si} has r instantiations, $\{o^{si}_1, ..., o^{si}_r\}$, where r is the product of the state number of each node in the group O^{si}. The following table gives the value and probability distribution for w_i:

w_i	$p(w_i\|\theta)$	$p(w_i\|\neg\theta)$
$ln\frac{p(o^{si}_1\|\theta)}{p(o^{si}_1\|\neg\theta)}$	$p(o^{si}_1\|\theta)$	$p(o^{si}_1\|\neg\theta)$
...	...	...
$ln\frac{p(o^{si}_r\|\theta)}{p(o^{si}_r\|\neg\theta)}$	$p(o^{si}_r\|\theta)$	$p(o^{si}_r\|\neg\theta)$

(13)

Based on the table, the expected value μ, and the variance σ for each w_i can be computed as:

$$\mu(w_i|\theta) = \sum_{j=1}^{r} p(o^{si}_j|\theta) ln\frac{p(o^{si}_j|\theta)}{p(o^{si}_j|\neg\theta)} \tag{14}$$

$$\sigma^2(w_i|\theta) = \sum_{j=1}^{r} p(o^{si}_j|\theta) ln^2\frac{p(o^{si}_j|\theta)}{p(o^{si}_j|\neg\theta)} - \mu^2(w_i|\theta) \tag{15}$$

By the central limit theorem, the expected value and the variance of W are:

$$\mu(W|\theta) = \sum_{i=1}^{g} \mu(w_i|\theta) \tag{16}$$

$$\sigma^2(W|\theta) = \sum_{i=1}^{g} \sigma^2(w_i|\theta) \tag{17}$$

From Equation 10, for $W \sim N(\mu(W), \sigma^2(W))$, $\phi \sim LogN(\mu(W)+c, \sigma^2(W))$, where LogN denotes Log normal distribution. Thus,

$$\begin{aligned} & p(\phi^*_{kl} \le \phi < \phi^*_{ku}|\theta) \\ = & \frac{1}{\sigma(W|\theta)\sqrt{2\pi}x} \int_{\phi^*_{kl}}^{\phi^*_{ku}} e^{\frac{-(\ln x - \mu(W|\theta)-c)^2}{2\sigma^2(W|\theta)}} dx \end{aligned} \quad (18)$$

$p(\phi^*_{kl} \le \phi < \phi^*_{ku}|\neg\theta)$ can be computed in the same way by replacing θ with $\neg\theta$ in the previous equations.

By combining Equation 1, 4, 11, and 18, VOI can be computed efficiently.

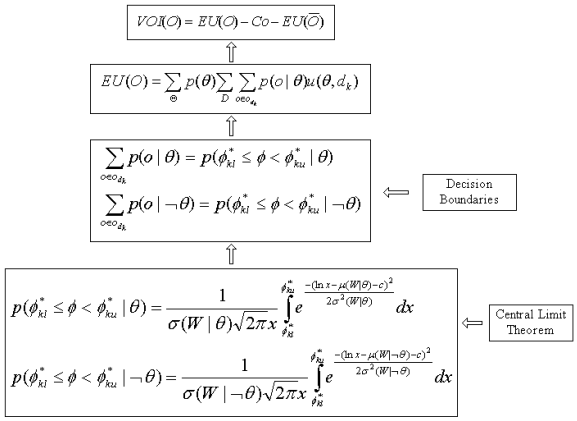

Figure 2: The key equations to compute VOI

As shown in Figure 2, to compute $VOI(O)$ efficiently, the key is to compute $EU(O)$, which leads to approximate $\sum_{o \in o_{d_k}} p(o|\theta)$ with Log normal distributions by exploiting the central limit theorem and decision boundaries.

Generalization

In the previous algorithm, the node Θ only allows two states although there is no such a limitation for the other nodes. In this subsection, we extend the algorithm to the case that Θ can have several states. Assume Θ has h states, θ_1, ..., θ_h, and still, d has q rules, $d_1, ..., d_q$, similarly to Equation 10,

$$\begin{aligned} & \frac{p(\theta_i|O)}{p(\theta_h|O)} = \frac{p(O^{s1}|\theta_i)}{p(O^{s1}|\theta_h)} ... \frac{p(O^{sg}|\theta_i)}{p(O^{sg}|\theta_h)} \frac{p(\theta_i)}{p(\theta_h)}, i \ne h \\ & \Rightarrow ln\frac{p(\theta_i|O)}{p(\theta_h|O)} = \sum_{k=1}^{g} ln\frac{p(O^{sk}|\theta_i)}{p(O^{sk}|\theta_h)} + ln\frac{p(\theta_i)}{p(\theta_h)} \\ & \Rightarrow ln\phi_i = \sum_{k=1}^{g} w^i_k + c_i, \end{aligned}$$

$$\text{where } \phi_i = \frac{p(\theta_i|O)}{p(\theta_h|O)},\ w^i_k = ln\frac{p(O^{sk}|\theta_i)}{p(O^{sk}|\theta_h)},\ c_i = \ln\frac{p(\theta_i)}{p(\theta_h)} \quad (19)$$

Let $W_i = \sum_{k=1}^{g} w^i_k, i \ne h$, W_i still has a Gaussian distribution. The similar method in the previous subsections can be used to compute the variance and mean. For the new defined w^i_k, Table 13 can be modified as follows:

w^i_k	$p(w^i_k\|\theta_1)$	...	$p(w^i_k\|\theta_h)$
$ln\frac{p(o^{sk}_1\|\theta_i)}{p(o^{sk}_1\|\theta_h)}$	$p(o^{sk}_1\|\theta_1)$	...	$p(o^{sk}_1\|\theta_h)$
...	...	...	...
$ln\frac{p(o^{sk}_r\|\theta_i)}{p(o^{sk}_r\|\theta_h)}$	$p(o^{sk}_r\|\theta_1)$	...	$p(o^{sk}_r\|\theta_h)$

(20)

Thus,

$$\mu(w^i_k|\theta_j) = \sum_{l=1}^{r} p(o^{sk}_l|\theta_j) ln\frac{p(o^{sk}_l|\theta_i)}{p(o^{sk}_l|\theta_h)} \quad (21)$$

$$\sigma^2(w^i_k|\theta_j) = \sum_{l=1}^{r} p(o^{sk}_l|\theta_j) ln^2\frac{p(o^{sk}_l|\theta_i)}{p(o^{sk}_l|\theta_h)} - \mu^2(w^i_k|\theta_j) \quad (22)$$

Similar to Equation 16, the expected value and the variance of W_i can be obtained as follows:

$$\mu(W_i|\theta_j) = \sum_{k=1}^{g} \mu(w^i_k|\theta_j), 1 \le i \le h-1, 1 \le j \le h \quad (23)$$

$$\sigma^2(W_i|\theta_j) = \sum_{k=1}^{g} \sigma^2(w^i_k|\theta_j) \quad (24)$$

Let $f_{\theta_j}(\phi_i)$ denote the probability density function of ϕ_i given θ_j, therefore, $f_{\theta_j}(\phi_i) \sim logN(\mu(W_i|\theta_j), \sigma^2(W_i|\theta_j))$.

Even though $f_{\theta_j}(\phi_i)$ can be easily obtained, it is still necessary to get the decision boundaries for each optimal decision in order to efficiently compute $\sum_{o \in o_{d_k}} p(o|\theta_j)$. A set of linear inequality functions need to be solved. For example, if d_k is the optimal action, $Eu(d_k)$ must be larger than or equal to the expected utility of taking any other action; through which, a set of linear inequality functions can be obtained:

$$\begin{aligned} & p(\theta_1)u_{1k} + p(\theta_2)u_{2k} + ... + p(\theta_h)u_{hk} \\ & \ge p(\theta_1)u_{1j} + ... + p(\theta_h)u_{hj}, \forall j, 1 \le j \le q, j \ne k \\ \Rightarrow\ & p(\theta_1)u_{1k} + ... + (1 - p(\theta_1) - ... - p(\theta_{h-1}))u_{hk} \\ & \ge p(\theta_1)u_{1j} + ... + (1 - p(\theta_1) - ... - p(\theta_{h-1}))u_{hj} \\ \Rightarrow\ & \frac{u_{1k}-u_{1j}}{u_{hj}-u_{hk}} \cdot \frac{p(\theta_1)}{p(\theta_h)} + ... + \frac{u_{(h-1)k}-u_{(h-1)j}}{u_{hj}-u_{hk}} \cdot \frac{p(\theta_{h-1})}{p(\theta_h)} \ge 1 \end{aligned} \quad (25)$$

We assume $u_{hj} - u_{hk} > 0$; otherwise, just change "$\ge$" to "$\le$" in the last equation.

Let A_k be the solution region of the linear inequality system, then

$$\sum_{o \in o_{d_k}} p(o|\theta_j) = \int_{A_k} f_{\theta_j}(\phi_1)...f_{\theta_j}(\phi_{h-1})dA_k \quad (26)$$

From this, $VOI(O)$ can be computed by combining Equation 1, 4, and 26.

Now, we analyze the computational complexity of the proposed approximation algorithm compared to the exact computation method. For simplicity, assume the state number of each O_i node is m and there are n nodes in the set O. Then the computational complexity of the exact VOI computation method is $O(hm^n)$, where h is the state number of the Θ node. With the approximation algorithm, the computational complexity is reduced to $O(qhm^k)$, where q is the number of decision rules for the decision node D, and k is the number of O_i nodes in the maximum group

among $\{O^{s1}, ..., O^{sg}\}$. In the best case, if all the O_i nodes are conditionally independent given Θ, the time complexity is just $O(qhm)$, which is linear. In the worst case, if all the O_i nodes are dependent, the time complexity is $O(qhm^n)$. However, usually, in most real-world applications, k will be much less than n, thus the approximate algorithm tends to be very efficient, as will be shown in the experiments.

Active Sensory Action Selection

Till now, we have presented the approximation algorithm for efficiently computing non-myopic VOI. Let $A = \{A_1, ..., A_m\}$ be a collection of sets over O. To get an optimal sensor set, we can enumerate all the sensor sets in A and compare their VOIs. However, it is impractical in this way since the size of A is increasing exponentially as the number of sensors. In practice, we use the algorithm as follows:

Algorithm 1: SensorSelect(ID,m)

Let $G = \{A_i : |A_i| = m, A_i \in A\}$;
$A_1^* \leftarrow \arg\max\{VOI(A_i) : |A_i| < m, A_i \in A\}$;
$A_2^* \leftarrow \emptyset$;
for each $A_i \in G$ **do**
  $G' \leftarrow O \setminus A_i$; $max \leftarrow VOI(A_i)$;
  while $G' \neq \emptyset$ **do**
    $O^* \leftarrow \arg\max\{VOI(O_i \cup A_i) : O_i \in G'\}$;
    if $VOI(O^* \cup A_i) > max$ **then**
      $A_i \leftarrow O^* \cup A_i$; $max \leftarrow VOI(A_i)$;
    $G' \leftarrow G' \backslash O^*$;
  if $VOI(A_i) > VOI(A_2^*)$ **then** $A_2^* \leftarrow A_i$;
Return $A^* = \arg\max(VOI(A_1^*), VOI(A_2^*))$;

As shown in the above pseudo code, on the first phase, Algorithm 1 gets a solution A_1^* by enumerating all possible l-element ($l < m$) subsets for some constant m. On the second phase, Algorithm 1 starts from all possible m-element subsets and then uses a greedy approach to supplement these sets so that to get a solution A_2^*. Finally, the algorithm outputs A_1^* if $VOI(A_1^*) > VOI(A_2^*)$ and A_2^* otherwise. The time complexity of the algorithm is $O(n^{m+1} \log n)$. It also requires to compute VOI for the sets whose size could be as large as n. If without the VOI approximation algorithm, even the greedy approach with m=1 could be very slow due to the computation of VOI. Fortunately, the proposed VOI approximation algorithm can significantly speed up the selection procedure.

Experiments

We perform experiments to demonstrate the performance of the VOI approximation algorithm and the sensor selection algorithm in the proposed active fusion framework respectively. Each ID is generated with random structure (but satisfies the general structure of Figure 1) and parameters, whose maximal number of nodes is 50. The number of sensors is limited to 11 due to the exponential computational time behind the brute-force approach to get the ground-truth.

Efficient Non-myopic VOI Computation

50 different IDs are generated. Each ID is parameterized with 20 sets of different conditional probability tables and utility functions, which yields 1000 test cases. For each test case, the VOIs for different O subsets with the size from 3 to 11 are computed. Among half of the 1000 test cases, the Θ node has binary states, and in the other half Θ has three states. All the other random nodes and decision nodes have four states.

Figure 3 illustrates the results of the 1000 test cases. Chart(a) shows the average error rates while chart (c) shows the VOIs for one specific case, where the error rate is defined as the the difference of the approximated VOI and exact VOI divided by the exact VOI. As the size of the O set increases, the error rate decreases. We could run several much larger IDs with much more O nodes, and the curve would be more and more closer to zero. Here, we intend to show the trend and the capability of this algorithm. Chart(b) shows the average computational time with the exact computation vs approximation computation. When the size of O node is small, the computational time is similar. However, as the size becomes larger, the computational time of exact computation increases exponentially, while the computational time of the approximation algorithm increases almost linearly. Thus, the larger the O set size is, the more time the approximation algorithm can save. If the state number of each O_i node is larger than 4, the saving would be more obvious.

Active Sensor Selection

The same 1000 testing cases are used for testing the sensor selection algorithm. The ground-truth of the optimal sensor subset is obtained by a brute-force approach. The value of m in Algorithm 1 is set as 3 in the experiments. Overall, the average VOI rate (the ratio between the VOI of the selected sensor set and that of the ground-truth optimal sensor set) is 0.97, and the running time rate (the ratio between the computational time with the proposed sensor selection algorithm and that with the brute-force method) is 0.02. Figure 3(d) displays the sensor selection results in 50 testing cases. Obviously, with a random selection method, the VOI of the selected sensor set is much less than the VOI of the optimal sensor set in most cases, while the proposed sensor selection can always select a good sensor set.

An Illustrative Application

We use a real-world application to demonstrate the advantages of the proposed framework and the algorithms. Figure 4(a) is a reflection of the proposed active fusion model for stress recognition and user assistance in a real-time system. The system decides an optimal sensory action set to collect evidence with the proposed sensor selection algorithm. The collected evidence is then propagated through the model and the posterior probabilities of user stress is computed with the dynamic inference technique. In the meanwhile, the system determines the optimal assistance that maximizes the overall expected utility.

Figure 4(b)(c) shows the experimental results of the VOI approximation algorithm for the stress model. We enumerate all the possible sensor sets and then compute the VOI for each set. Chart(b) illustrates the average VOI error rates for different sensor sets with the same size and chart (c) displays the computational time. Similarly to the simulation experiments, the error decreases as the size of O set increases, and

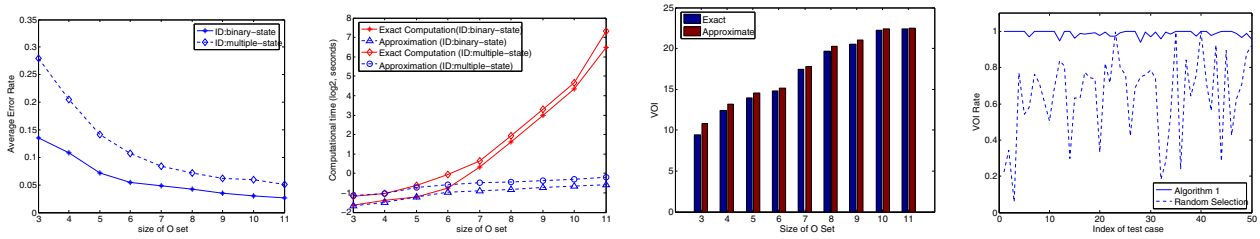

(a) (b) (c) (d)

Figure 3: Results for the 1000 test cases: (a) average error rates with the VOI approximation algorithm; (b) computational time (log(t), unit is second); (c) VOIt vs VOI for one test case; (d) performance of the proposed sensor selection algorithm vs random selection.

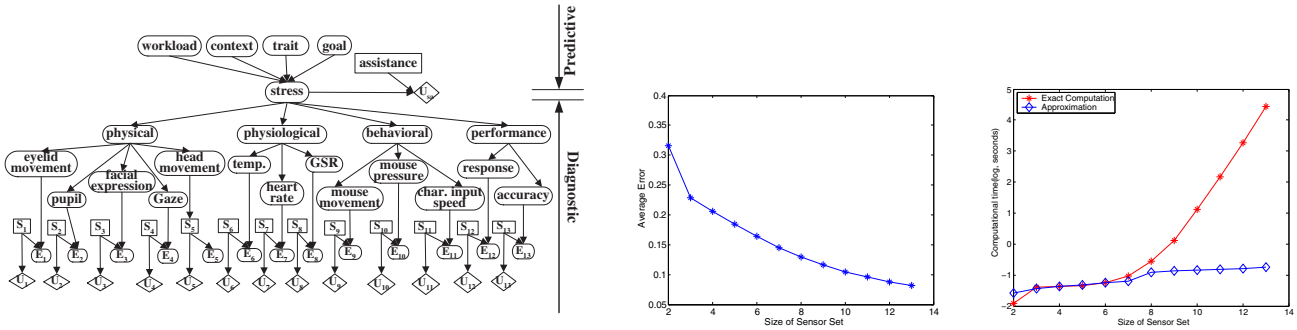

(a) (b) (c)

Figure 4: (a) An Influence Diagram for recognizing human stress and providing user assistance. More detail please refer to (Liao *et al.* 2005); (b)VOI error rate; (c)computational time(log(t), unit is second).

the computational time increases almost linearly with the VOI approximation algorithm. Furthermore, based on the approximated VOIs, the sensor selection algorithm can always return an optimal or near-optimal sensor set efficiently.

Conclusions

We adapt an Influence Diagram as an active fusion model for representation, integration, and inference of uncertain sensory information of different modalities at multiple levels of abstraction. Based on this framework, we present an approximation algorithm to compute non-myopic VOI of sensor sets efficiently by exploring the central-limit theory. Such an algorithm, together with the improved greedy approach, make it feasible to choose an optimal sensor set with significantly reduced time complexity. Both the experiments with synthetic data and real data from a real-world application demonstrate the proposed framework together with the proposed algorithms are able to efficiently and effectively select the optimal sensory actions for timely decision-making.

Acknowledgments

The work described in this paper is supported in part by a grant from AFOSR with grant number F49620-03-1-0160.

References

Dittmer, S., and Jensen, F. 1997. Myopic value of information in influence diagrams. *UAI97* 142–149.

Ertin, E.; Fisher, J.; and Potter, L. 2004. Maximum mutual information principle for dynamic sensor query problems. *IPSN'03* 405–416.

Fassinut-Mombot, B., and Choquel, J.-B. 2004. A new probabilistic and entropy fusion approach for management of information sources. *Information Fusion* 5:35–47.

Heckerman, D.; Horvitz, E.; and Middleton, B. 1993. An approximate nonmyopic computation for value of information. *PAMI* 15(3):292–298.

Jensen, F. 2001. *Bayesian Networks and Decision Graphs*. New York: Springer Verlag.

Kalandros, M.; Pao, L.; and Ho, Y.-C. 1999. Randomization and super-heuristics in choosing sensor sets for target tracking applications. *Proceedings of the 38th IEEE Conference on Decision and Control* 2:1803–1808.

Liao, W.; Zhang, W.; Zhu, Z.; and Ji, Q. 2005. A decision theoretic model for stress recognition and user assitance. *Twentieth National Conference on Artifical Intelligence (AAAI)* 529–534.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann Publishers.

Pinz, A.; Prantl, M.; Ganster, H.; and Borotschnig, H. K. 1996. Active fusion - a new method applied to remote sensing image interpretation. *Pattern Recognition Letters* 17:1349–1359.

Poh, K. L., and Horvitz, E. 1996. A graph-theoretic analysis of information value. *UAI-96*.

Shachter, R. 1999. Efficient value of information computation. *UAI-99* 594–601.

Wang, H.; Yao, K.; Pottie, G.; and Estrin, D. 2004. Entropy-based sensor selection heuristic for target localization. *IPSN'04* 36–45.

Zhang, N. L.; Qi, R.; and Poole, D. 1993. Incremental computation of the value of perfect information in stepwise-decomposable influence diagrams. In *UAI-93*.