

Active Affective State Detection and User Assistance

Xiangyang Li

Department of Computer Science
University of Nevada, Reno, NV 89557
xyli@ieee.org

Qiang Ji

Electrical, Computer, and Systems Engineering Dept.
Rensselaer Polytechnic Institute, Troy, NY 12180
qji@ecse.rpi.edu

Abstract.

Intelligent user assistance systems face challenges of incomplete, uncertain and multiple modality sensory observations, user's changing internal state, and constraints in making decisions. We introduce a probabilistic framework to dynamically model user's affective state in such systems. A systematic mechanism performs purposive and sufficing information integration to infer user's affective state and provide correct assistance. We aim to actively infer the user's status and engage in appropriate assistance in a timely and efficient manner.

1. Introduction

Intelligent assistance systems are an important application area of user modeling, especially with the rapid development of pervasive and ubiquitous computing. For example, in every year many people are injured in car accidents because drivers are in dangerous status including fatigue, nervousness, or confusion. If we could distinguish these dangerous states in a timely manner, and provide assistance in terms of appropriate alerts, we may prevent many accidents from happening. However such systems face several challenges: 1) sensory data are often incomplete, uncertain, and from sources of different modalities; 2) sensory data are often dynamic and evolving over time to reflect change in the user's state; and 3) decisions about the user's need and the assistance must be rendered appropriately and in a timely and efficient manner under various constraints on time and resources.

We introduce a probabilistic framework based on the Dynamic Bayesian Networks (DBNs) and information theory to simultaneously address the above challenges. Firstly, a generic hierarchical probabilistic framework for user modeling is introduced to model the sensory observations, and the profile and contextual information related to the user's mental state. Secondly, this framework dynamically evolves and grows to account for temporal change in sensory observations as a result of the change in user's internal state. The DBNs allow the temporal information to be systematically incorporated via temporal causality. Thirdly, the proposed framework provides a mechanism that performs purposive and sufficing information integration in order to determine the user's status.

Specifically, instead of passively fusing the information that is available, this system first formulates an initial hypothesis about the user's current internal state and then actively selects the most informative sensory/questioning strategy in order to quickly and economically confirm or refute the hypothesized internal state. All these methods help the system to actively infer the user's need/state under uncertainty over time and engage appropriate assistance in a timely and efficient manner.

2. Bayesian Networks in User Modeling

Recently, there has been a significant surge in using BNs for plan recognition, user need inference, and affective state assessment. Plans encode a user's intentions and desires. Huber et al (1994) provide a uniform procedure for converting plans represented in a flexible procedural language to probabilistic belief networks. Pynadath and Wellman (1995) present a Bayesian framework describing the context, the mental state and planning process of an agent, and the consequences of the agent's actions.

Intelligent user assistance systems need the ability to adaptively accommodate user's specific need. The READY system (Bohnenberger et al, 2002) uses DBNs in a dialog system to adjust the policy in providing instructions, based on the recognized time pressure and cognitive load from the user's utterances, realized by a rule base that maps detected situations into actions. Microsoft's Lumiere project (Horvitz et al, 1998) is intended to help computer users with interactive interfaces by identifying their needs. DeepListener and Receptionist (e.g., Horvitz & Paek, 2000) augment the speech recognition in clarification dialogs by inferring user intentions associated with utterances. Costs and benefits can be calculated for different available actions and the action with the highest utility is executed. All such systems, however, do not distinguish actions and sensory tests and relies heavily on immersive interaction.

Affective computing (Picard, 1997) uses pattern recognition and information retrieval technologies for affective state assessment. Existing techniques include, discriminant analysis, fuzzy rules, neural networks, Bayesian learning, HMM model, Bayesian networks, etc. One group of them uses sensory




Figure 1 “Context-Profile-State-Observation” model, where self-pointing arrows indicate temporal links.

measures as predictor variables and applies classification algorithms without prior and context knowledge. The other group, represented by the Bayesian network and HMM models, represents the prior knowledge and expertise in graphic network form. They maintain the balance between the global and local representations and provide capabilities in handling the uncertainty and incompleteness in practical systems with the aid of the causal and uncertainty representation structure. Ball and Breese (2000) use a dynamic Bayesian network to assess the user’s affective state in terms of valence and arousal, using facial and speech information. Conati (2002) provides a dynamic Bayesian network model for assessing students’ emotion in educational games. The emotion states are modeled as consequences of how the current situation (action and help) fits with the student’s goals and preferences. Some body expressions and sensors are also used as evidence.

Finally, another area of research that is related to the proposed active sensing is fault detection and troubleshooting (Langseth & Jensen, 2003). It focuses on globally seeking the best action sequence for a problem setting where the actions and questions are not repeated. Their definitions of efficiency, cost, and value of information are useful for our task to efficiently and accurately infer user state.

In conclusion, the researchers have realized the benefits of DBNs and utility theory, and have begun to apply them to user assistance. Current research in these areas, however, is limited to passive inference, mostly affect-insensitive, and in a static domain. Compared with the Bayesian network systems discussed above, our system currently aims at two

objectives. First, non-intrusive and active user state inference. Our target is to design silent agents that use the most reliable and non-intrusive evidences, to provide the user with accurate and active assistance in a pervasive and ubiquitous computing environment. Second, dynamic and active sensor selection. The selection of sensor or sensors in such a system should not be done once and then forgotten, but needs to be continually and dynamically reevaluated. We focus more on refining sensors/questions dynamically using a local optimal strategy.

3. A Generic User Modeling Framework

Our generic framework to apply Bayesian networks to user modeling is the “Context-Affective State-Profile-Observation” model. It is used to infer the user’s affective state from their visual observations. As in **Figure 1**, such model captures the user’s profile, affective state, and the contextual information.

- Contextual component. The represents information about the specific environmental factors that can influence the user's affective state.
- Affective state component. This component represents the user’s emotional status. Typical affective states include fatigue, confused, frustration, fear, sad, and anger.
- Profile component. This models user’s ability and competitiveness in finishing the operations. This provides the adaptation capability of the model to individual users.
- Observation component. It is sensory observation of different modalities about user behaviors.

The affective state of the user and hidden nodes of the user’s visual, audio and behavioral status in

current time slice are influenced by the corresponding variables in the most recent time slice. The user profile could also have temporal links between time slices. However in this figure, we consider it unchangeable in a running session. This figure also outlines the causal relations between context, profile, state, and observation variables. The context and profile variables influence the user's state. The user's states lead to the evolution of visual, audio, and behavioral expressions.

4. Active Affective State Inference

Since we are often constrained by the time and resource we could use, and the strict requirement on the accuracy of assistance, purposive and sufficing information collection and integration are needed to infer about the user's affective state in a timely and economic manner. **Figure 2** shows a general view of the active user state detection system. We are interested in how to dynamically control (select actions and make decisions) the system that has a repertoire of sensors such that the system operates in a purposive manner. We selectively collect the observations of sensory variables, or even further reduce the uncertainty by asking user questions.

4.1 Active User State Inference

Mathematically, the user affective state inference problem may be viewed as a hypothesis detection problem, with hypothesis, $H = \{h_1, h_2, \dots, h_n\}$, representing the possible user state. The sensory observations $E = \{E_1, E_2, \dots, E_m\}$, have m diverse sensors. The goal is to estimate a posterior probability that $H = h_i$ is true given E , i.e., $P(H = h_i | E)$. According to the Shannon's measure of entropy, the entropy of a distribution over hypothesis H_i , given the hypothesis distribution in last time slice H_{t-1} is:

$$ENT(H) = -\sum_{h_{t-1}} p(h_t | h_{t-1}) \log p(h_t | h_{t-1})$$

The benefit of certain evidence can be measured by its potential to reduce the uncertainty with the hypothesis, called mutual information, i.e., the

differential entropy between prior to and after the sensory action. Given sensor E_i with a set of states (e_1, e_2, \dots), this mutual information can be denoted as:

$$\begin{aligned} I(E_i) &= ENT(H) - \sum_{e_i} p(e_i) ENT(H | e_i) \\ &= -\sum_{h_{t-1}} p(h_t | h_{t-1}) \log p(h_t | h_{t-1}) \\ &\quad + \sum_{e_i} [p(e_i) \sum_{h_t} p(h_t | h_{t-1}, e_i) \log p(h_t | h_{t-1}, e_i)] \end{aligned}$$

This formula is the fundamental equation for computing uncertainty reducing potential for H due to E . We could easily extend it to consider the case of multiple sensors, $E = \{E_1, \dots, E_n\} \subseteq E$. The probabilities in the above equation are readily available from the forward and backward inference propagation based on hypothesis beliefs for last time slice. For example, $p(h_t | h_{t-1}, e_i)$ is the posterior probability of hypothesis for current time slice given the a state combination on sensor variables.

Acquiring information incurs cost. The cost may include the cost of information retrieval, the time to include the information from source into the fusion system, computation time for sensory data processing, and hardware execution time. We consider the sensor cost C of selecting E_i , where all costs are assumed to be incorporated into the same equivalence, using the following formula:

$$C(E_i) = \frac{C_i}{\sum_{j=1}^m C_j}$$

where C_i is the cost to acquire the information from sensor i . Combining uncertainty reducing potential and information acquisition cost, we form the expected utility given sensor E_i as:

$$EU(E_i) = \alpha I(E_i) + (1 - \alpha)C(E_i)$$

where α is the balance coefficient between the two terms. The optimal sensor action can be found by using the following decision rule:

$$E^* = \arg \max_{E_i} EU(E_i)$$

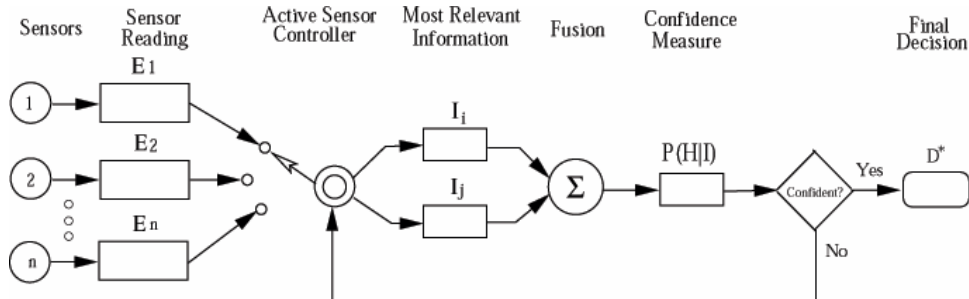


Figure 2 Active state assessment system overview.

4.2 Some Issues in Active Inference

Here we discuss some complicated situations in sensor selection: inference degradation, selection of multiple sensors, and multiple hypotheses.

One problem we encounter is inference degradation. The “degradation” here is defined as the compromise on inference ability when a small set of sensors are repeatedly selected. Due to the mutual information definition, a sensor tends to yield a higher value in calculation, if its conditional probability table has very uneven entries for configurations of parent states. It is also observed from the implementation. This is reasonable since our principal objective in active fusion is to seek the sensor which could most effectively distinguish various affects of the subject. However, on the other hand, the situation that one or a very small set of sensors dominate the selection depresses the advantage of fusing multimodal information for better accuracy and robustness. Furthermore, such superiority of these sensors in mutual information comes partially from using the expected potential. In such calculation, all possible states of sensor are inspected and the whole network structure impacts the result. But later on, only one state is instantiated for the selected sensor. In this sense, such criterion has bias in selection process. To reduce inference degradation, mutual information should not be the only factor in determining sensor’s benefit. An obvious solution is to force more sensors into engagement. In our study we use a taboo list to exclude most recently selected sensors from current selection decision, taking the history information in consideration. A list length of 1 means this list contains the sensor selected most recently.

The second issue is efficient multiple sensor selection when more than one sensors could be opened together. We have to consider the mutual information for every possible combination of sensors in seeking the optimal configuration. This is a typical NP-hard problem, demanding intense computation with the increase of sensors. Seeking an affordable search algorithm for this case is a challenge for future research. In our research, we use a greedy or myopia strategy that calculates the utility for each sensor and ranks them accordingly. Then we choose the number of sensors needed from the top.

We also consider the case where several interesting affects exist in one model. There are different views with regard to whether different affective states could coexist. Here we use multiple binary affect nodes since this approach could accommodate exclusive affects too. A constraint could be put on these nodes to keep the relation of the positive states among them. Let $I(H_j, E_i)$ be the mutual information of E_i to

hypothesis j . We rewrite the sensor’s mutual information.

$$I'(E_i) = \sum_j w_j I(H_j, E_i)$$

where w_j is the weight to distinguish the importance of affects. In our experimentation, we set this weight as the current belief for the hypothesis’s positive state, i.e., $w_j = p(H_j = \text{positive})$. This means we give the suspicious hypothesis more importance.

5. Decision on Assistance

There are two key questions to answer: when we should provide the help and what help we should provide. The first question normally requires control thresholds on the probability distribution for state variables. We calculate a State Level (SL) on the probability distribution of affective states:

$$SL = \sum_{h_i} w(h_i) p(h_i)$$

where $w(h_i)$ is the weight for the i th state of an affect, indicating this state’s seriousness level. For example, this weight could be set as zero for the negative state of a dangerous affect such as fatigue. Then we could set an Engaging Threshold (ET) on SL . If SL is greater than ET , we provide assistance for users. Moreover, we may want to use a set of state levels when there are different affective state variables that are not exclusive with each other. In practice, a SL smoothed over time is more appropriate as a reliable indicator. In our experimentation, the SL s are smoothed over three time slices.

What assistance to provide depends on user’s current state and the utility of assistance. The utility of assistance represents the optimal trade-off between its benefits and its cost. The benefits focus on the beneficial consequence of the assistance. One measure of the benefit could be the assistance’s potential to return the user from an anomaly state to his/her nominal state. The benefit could be calculated by assessing the cross product of the situations and these assistances, through psychological experiments on a population of users, or some assessment tools like using unidimensional or multidimensional scaling. The cost includes the computational cost, the potential of annoying the user, the physical cost, and the cost of not providing or delaying the assistance.

The utility of assistance is also impacted by the user’s current status, including the affective state, the current task goal, the cause, and the user’s tolerance to assistance, shown in **Figure 3**. “Task” shows the user’s current interest, such as choosing some icon or button. “Cause” is the explanation of reasons for the subject’s state. “Tolerance” is like a switch control variable determining the intervention degree the user

would agree on. A utility form considering all these factors needs much more research effort. In this paper, we just give the simplest utility calculation based on the beliefs of the user’s affective state.




Figure 3 Active assistance model.

Let A_j represent the j th assistance in consideration; let $G_B(A_j, h_i)$ and $G_C(A_j, h_i)$ represent the benefit and cost for assistance j respectively, given user’s current state h_i . Then the expected utility of assistance A_j may be defined as:

$$EU(A_j) = \sum_i G_B(A_j, h_i) p(h_i) - \sum_i G_C(A_j, h_i) p(h_i)$$

And then the best assistance is determined via:

$$A^* = \arg \max_{A_j} EU(A_j)$$

6. Evaluation Model

We use subjective parameters and simulated data to evaluate this framework. The task is to detect whether a computer operator is among “fatigue,” “nervousness” and “confusion” mental affects, using visual cues about facial expression, eyelid, gaze, and “query.” The implementation is in MATLAB using the BNT toolkit. The inference algorithm uses the junction tree engine. Description of these discrete variables is in **Table 1**. We use three separate nodes for the mental affects because we do not stipulate that these states are exclusive from each other. Explanation for these visual cues and “query” is given in **Table 2**.

The network parameters include prior probabilities for “context” and “profile” nodes, and conditional probabilities for the links. In our experimentation, the prior probabilities are all set even, 0.5 for each state since the root nodes are all binary.

The conditional probabilities of mental affects between two consecutive time slices are also called “transitional” probabilities. Normally this probability between same states, e.g., “positive” to “positive”, is high, if we consider a user’s mental state remain relatively stable. Transient affects such as “confusion” may come and go more quickly, with lower such probability. Transitional probability

between opposite states, e.g., “positive” to “negative,” is much lower correspondingly. In experimentation, the transitional probability between the same states of “fatigue” is 0.9, while 0.85 for two other affects.

Table 1 Variables used in the evaluation model.

Component	Variables	States
Context	Context	complex/simple
	Physical condition	strong/weak
Profile	Skill	strong/weak
	Fatigue	negative/positive
Mental state	Nervousness	negative/positive
	Confusion	negative/positive
	Facial expression	neutral, happiness, sadness, anger, surprise, disgust, fear
Sensory observation	Eyelid PERCLOS	high/normal/low
	Eyelid AECS	fast/normal/slow
	Gaze spatial distribution	high/even/low
	Gaze fixation ratio	high/even/low
	Query	fatigued/confused/no

Table 2 Explanation of visual cues and “query”.

Facial expression	Expressions from the FACS system.(11)
Eyelid PERCLOS	Percentage of eyelid closure over the pupil over time. (12)
Eyelid AECS	Average eye closure/open speed over time. (13)
Gaze spatial distribution	The ratio of gaze staying outside the computer screen in the overall time. (14)
Gaze fixation ratio	The ratio of gaze fixation over time. (15)
Query	Answers to direct questioning in the form “are you __?” among “fatigued, confused, comfortable.” (16)

Other conditional probability values are from a survey. We focus more on the working mechanism than a fidelity model. However we have too many conditional probabilities in terms of all combinations of states in parent and child nodes. Instead in our survey we first investigate the conditional probability on a single parent, e.g., the probability of gaze fixation ratio being high given the subject is fatigued. Noisy-or assumption or the extensions have been used to reduce the number of probabilities to estimate from these one-on-one probabilities. While such extension could deal with arbitrary input and output nodes, the physical meaning behind the nary variables and the output function are very opaque to represent and interpret. Thus we approximate the

combinatorial conditional probabilities as:

$$p(C | A_1, \dots, A_n) \approx \prod_i P(C | A_i)$$

where A_1, \dots, A_n are parents of C . The assumption is that these affect function independently to raise the external expressions in visual sensory channels.

Table 3 Six settings in evaluation experimentation.

Setting	Notes
Passive I	Randomly select 1 sensor.
Active I	Actively select 1 sensor.
Active II	Actively select 1 sensor with Taboo List.
Passive II	Randomly select two sensors.
Active III	Actively select two sensors.
Assistance	Provide assistance in active fusion III.

Different sensor activation and assistance strategies are listed in **Table 3**. In assistance setting, all ETs are set to 0.8. “Query” provides more accurate measure of the user’s mental state. But it is very intrusive to the subject. We demonstrate the function of query by only turning on it once as the last confirmation before any assistance. In all active fusion settings, the sensor costs are all set as zero, i.e., α is 1.

Different scenarios are used, including three “affective” scenarios (“fatigue,” “nervousness,” “confusion”), and “normal” scenario. At the beginning of each scenario, the beliefs for positive and negative states of each affect variable are all set as 0.5. In each scenario, the selected sensor is instantiated using the corresponding sensor state for the current time slice. In determining these sensor states in these scenarios, we first do a forward propagation using the same Bayesian network, setting probabilities for the affect variables. For example, in the scenario where the subject is fatigued, the probability for the positive state of “fatigue” variable is set as 99%, while the probabilities of the positive states of other two affective state variables are set as 1%. After propagation, each sensor has a probability distribution associated with its states. Each sensor turns out the state with the highest probability. In the sense of reliability, such sensory channel could always catch the observation from the subject that is the most indicative signal for the underlying affect.

7. Experimental Result Analysis

The posterior probability for the positive state of each affect variable is recorded in each time slice, as well as the calculated SLs . Thereafter we call this probability the belief of the corresponding affect, e.g., the belief of “fatigue.” This belief is the measure to summarize the results about comparison of active fusion with passive fusion, sensor sequences, usage

of taboo list, and assistance process.

7.1 Active Inference versus Passive Inference

We first examine the “normal” scenario since we focus on “affective” scenarios in the rest of the section. **Figure 3** shows the belief curves in passive and active fusions respectively. As the curves show, active fusion settings (on the right) detect the underlying status of the subject more quickly. In “normal” scenario, all three probabilities drop below 0.5. Although the corresponding passive fusion settings could detect the same trends in these beliefs, they are not as efficient as active fusion.




Figure 4 The beliefs for passive and active inference.

Now we examine the belief difference between passive and active fusions in each time slice. In three “affective” scenarios, we focus only on the difference for the underlying affect, i.e., fatigue for “fatigue” scenario, and so forth. Ignoring other affects, we plot this difference against time slice for the settings using one and two sensors respectively, shown in **Figure 5**. In the graphs, the belief of active setting subtracting that of passive setting produces a set of points for each scenario. Most points lie above the X-axis, meaning that in most time slices, the uncertainty reduction for the underlying affect of each scenario is more significant in active fusion than in passive fusion. However, we notice there are some

exceptions though only occurring in a minority of all cases, such as in the “fatigue” scenario using one sensor. In these cases, the points lie below the X-axis, meaning the passive fusion gives higher beliefs for the underlying affect. And finally, the superiority of active setting to passive setting is more evident at the beginning time slices. Although in the late stage, passive settings likely have higher certainty in some cases. This is the possible what we call “inference degradation.”

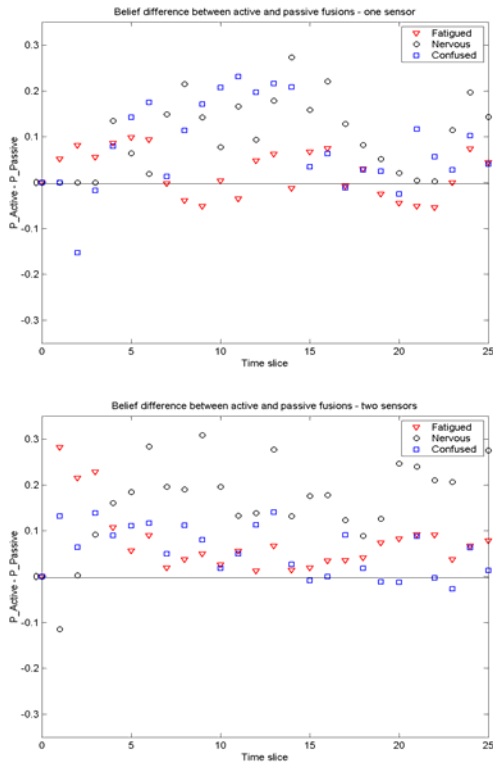


Figure 5 The belief difference ($P_{Active} - P_{Passive}$) of the underlying affect, using one and two sensors.

Table 4 The number of time slices needed to reach the target threshold on SL for three “affective” scenarios, where “n/a” indicates the threshold is never reached during the 25 time slices.

Scenario	One sensor (threshold = 0.8)		Two sensors (threshold = 0.9)	
	Passive	Active	Passive	Active
“Fatigue”	7	5	7	3
“Nervousness”	n/a	14	n/a	6
“Confusion”	16	6	11	4

We could also compare the uncertainty reduction abilities by setting a target threshold on the underlying affect belief, and compare in different settings the number of time slices used to first reach this threshold. This is a good measure since in

practice we could regard this threshold as the control threshold for assistance engagement. Moreover we would like to use the individual SL to be the substitute of affect belief because it provides a more reliable estimator. Here we set such target threshold as 0.8 for one sensor and 0.9 for two sensors in each scenario. We draw similar conclusion from **Table 4**.

7.2 Sensor Sequences in Active Fusion

We want to examine the sensor sequence selected in active fusion. Active fusion selects the sensors with the highest utility in each time slice. This utility may change along with time even in the same scenario. **Table 5** shows the sensor sequences for different scenarios without sensor costs. Because the way we assign the initial beliefs, the first sensor selected for all scenarios are all the same, “AECS.” Then with the change of affective hypothesis beliefs, different sensors may be selected, based on the merit of mutual information and sensor cost. However, we notice that not all sensors are selected. More specifically only sensors 12 (PERCLOS), 13 (AECS) and 15 (gaze fixation ratio), are ever selected in all scenarios. Also in the late time slices, the sensor sequence is fixed, with certain sensors repeated.

Table 5 Sensor sequences in each scenario in active fusion setting where one sensor is selected in each time slice.

Time	Fatigued	Nervous	Confused	Normal
1	13	13	13	13
2	13	13	12	12
3	13	13	12	15
...
13	13	13	12	15
14	13	13	12	12
...

7.3 Usage of Taboo List

When we recall the “inference degradation” in comparing passive and active fusion performances, now we see repetition of sensors in active fusion in sensor sequence analysis. We use the taboo list in active fusion to force more sensors into selection to try to counter this problem. In **Figure 6**, we compare the active fusion setting using taboo list of length 1 with the previous setting without the taboo list, opening one sensor in each time slice. Similarly, we plot the belief difference between the two settings of the underlying affect variable for the three “affective” scenarios. From the produced points, we observe that such taboo list improves the performance of “fatigue” and “nervousness” scenarios in late time slices. However, it lowers the hypothesis beliefs for “confusion” scenario from the beginning. Thus,

rather than we conclude the effect of such a way using utilization history of sensors, an advice of delicate deployment is more important for this method to alleviate the inference degradation.

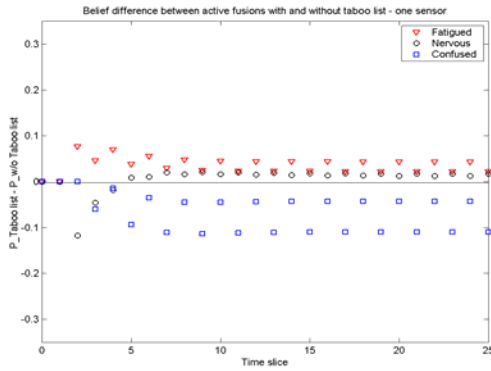


Figure 6 The belief difference of the underlying affect in each “affective” scenario, for active fusions with and without taboo list used, opening one sensor.

7.4 Assistance Process

The *SLs* and corresponding utilities for different assistances when the assistance is engaged in are given in **Table 6**. Although the individual *SL* for the underlying affect in each “affective” scenario reaches the *ET* (0.8 here) very fast, the assistance is available after the fifth time slice since we want to accumulate evidence for enough time. For each scenario, we calculate the utility for each assistance and choose the assistance with the highest utility. In our case, the assistance is appropriate in all scenarios, i.e., warning for “fatigue” scenario, emphasis for “nervousness” scenario for the subject to focus, and simplification of unrelated information for “confusion” scenario. While the assistance soothes the danger of one affect, it may aggravate other affects, e.g., the warning may intensify “confusion.” This tells us the importance of accurately detecting the subject’s status in order to provide appropriate assistance.

Table 6 *SLs* and assistance utilities in “affective” scenarios, with *ET* = 0.8 and opening two sensors.

Scenario	<i>SL</i>			Utility		
	F	N	C	W	E	S
F	0.87	0.53	0.19	0.60	0.23	-0.61
N	0.10	0.85	0.75	-0.38	0.68	-0.13
C	0.07	0.12	0.81	-0.43	0.09	0.52

*F-fatigued, N-nervous, C-confused, E-emphasis, W-warning, S-simplification

8. Conclusion

This research makes contributions in dynamically and systematically modeling the user affective state,

and performing active information fusion so that the user’s state and need can be determined and met in a timely and efficient manner. The active inference strategy provides better performance in evaluation experimentation. We also notice that such an affective state detection system alone could not fully fulfill very accurate assistance. Further research is ongoing in our lab to integrate multiple and heterogeneous models in such task.

Acknowledgements

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

9. References

- Ball, G., & Breese, J. (2000). Emotion and personality in a conversational agent. *Embodied Conversational Agents*. Cambridge, MA: MIT Press.
- Bohnenberger, T., Brandherm, B., Großmann-Hutter, B., Heckmann, D., & Wittig, F. (2002). Empirically grounded decision-theoretic adaptation to situation-dependent resource limitations. *Künstliche Intelligenz*, Vol. 3.
- Conati, C. (2002). Probabilistic assessment of user's emotions in educational games. *Applied Artificial Intelligence, special issue on “Merging Cognition and Affect in HCI”*, Vol. 16, No. 7-8.
- Horvitz, E., Breese, J., Heckerman, D., Hovel, D., & Rommelse, K. (1998). The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Pp. 256-265.
- Horvitz, E., & Paek, T. (2000). DeepListener: Harnessing expected utility to guide clarification dialog in spoken language systems. *6th International Conference on Spoken Language Processing (ICSLP 2000)*. Beijing.
- Huber, M. J., Durfee, E. H., & Wellman, M. P. (1994). The automated mapping of plans for plan recognition. *Workshop on Distributed Artificial Intelligence*, Pp. 137-152. Lake Quinalt, WA.
- Langseth, H., & Jensen, F. V. (2003). Decision theoretic troubleshooting of coherent systems. *Reliability Engineering and System Safety*, Vol. 80, No. 1, Pp. 49-61.
- Picard, R. (1997). *Affective Computing*. MIT Press, Cambridge, MA.
- Pynadath, D. V., & Wellman, M. P. (1995). Accounting for context in plan recognition, with application to traffic monitoring. *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, Pp. 472-481. Morgan Kaufmann, San Francisco.