Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Fuzzy Expert Systems For Sequential Pattern Recognition For Patient Status Monitoring in Operating Room

Joel Xue, Michael Krajnak, *GE Healthcare*

Abstract— **In this paper, we present several fuzzy inference systems for monitoring patient status in an operating room. The algorithms used include recursive fuzzy inference (RFIS), and non-recursive with sequential patterns as inputs. The RFIS algorithm combines current patient status data with previous output of the inference system, therefore is able to reinforce the current finding based on previous sequential system output. The results show that the RFIS system can be tuned towards higher sensitivity for more critical status, while generating a smoother inference output.**

I. INTRODUCTION

FUZZY logic and fuzzy inference systems (FIS) have been widely used in control and pattern recognition applications. There are quite a few applications of fuzzy logic in medical decision support [6]. A key advantage of using fuzzy logic in clinical decision support is that the rules can be programmed and easily understood by clinicians, unlike neural networks and other regression approaches, where the system behaves more like a black box [7].

Most fuzzy logic systems have a feed forward structure, which consists of an input layer with membership functions for each input parameter, a rule engine which combines input membership functions in a simple if-then grammar, and a defuzzification layer to combine the results from the rule engine [1]. This structure works well for static pattern recognition tasks. However, in the application of temporal pattern recognition, where classifications are made based on a sequence of input patterns, the feed-forward structure has its limitations. The example of sequential pattern recognition could be arrhythmia detection based on ECG and other real-time parameters, hypovolemia status monitoring, and anesthesia agent overdose monitoring in an operating room.

There could be two approaches to processing sequential patterns using a FIS: one method is to add sequential patterns to the input layer, so that the sequential elements of a pattern, x^0, x^{-1}, …, x^{-N} will each be treated as a separate input parameter with its own membership functions. The advantage of this method is that it does not change the structure of the FIS, and the disadvantage is the rapid increase of the number of rules due to increased input parameters. A second method is to take the previous output of a FIS and feed it back into the input layer to form a recursive FIS (RFIS) system.

A recursive model should provide a better indicator of a condition, e.g. have higher output values in an event region, due to reinforcement of the output from incorporating previous values. The feedback from the previous output should also result in the recursive FIS output rising later at the leading edge and falling slower at the trailing edge of the event than the non-recursive output. This is due to the fact that when the inference system considers prior history the output should be suppressed where there is no prior history of the condition and enhanced when the prior history indicates a high likelihood of an event. This effect should lead to a more accurate determination of the start and end times for a particular event. The construction and evaluation of an RFIS will be described in Section II.

We evaluated this concept by creating a baseline inference system for diagnosing hemodynamic crises brought on by anesthetic overdose. During anesthetic induction, surgical incision, or other periods where increased anesthetic affect is desired, agent levels are typically stepped up for some time period. If left up too long the increased agent can result in an undesirable decrease in blood pressure and heart rate. In a study of over 4,000 patients, 9% experienced severe hypotension in the period of 0-10 minutes after induction [5]. Intraoperative hypotension is generally associated with the post-induction phase of anesthesia and hypotension is a significant factor in long-term mortality and morbidity [3]. An automated system that detects this condition and notifies the clinician could help reduce the occurrence and severity of periods of low blood pressure.

Our goal is to verify the applicability of the RFIS model by using actual clinical data and experimenting with several expert system configurations in order to understand the effect of including the recursive input. We expect that adding a recursive input will increase the overall expert system performance without the additional complexity of adding sequential inputs in the input layer of the inference system. We compare several systems, a baseline without additional inputs, a modified system with a recursive input, and additional systems with time delayed inputs as additional features.

It was not our intent at this phase to produce a final set of clinical rules, but rather understand the behavior of several sequential input FIS configurations. We are considering the possibility of creating a high fidelity rule set using this technique for clinical use some time in the future.

II. MATERIAL AND METHODS

A. Data collection

For our experiments we used an annotated clinical data set collected as part of an ongoing research collaboration between GE Healthcare and the University of Michigan.

Manuscript received April 24 , 2006.

Joe Xue is with GE Healthcare Information Technology, Wauwautosa, WI 53226 USA (e-mail: joel.xue@med.ge.com).

Michael Krajnak is with GE Healthcare Information Technology, Milwaukee, WI 53201 USA (414-362-2495; e-mail: michael.krajnak@ med.ge.com).

This data set contains over 100 cases collected from a variety of high acuity surgical cases and includes a wide range of monitored physiological parameters. Each case has been annotated by two clinicians with the relevant events.

B.Feedforward FIS model

For our baseline we first created a three input FIS that uses heart rate trend, blood pressure trend, and expired anesthetic agent level to determine the degree of anesthetic overdose as shown in Figure 1. The initial inputs were selected based on clinician interviews regarding anesthetic overdose [4]. For the baseline configuration we limited the number of inputs to three. Each input is subject to set of filters designed to remove noise and signal artifacts. The rule set was tuned using a semi-automatic procedure [2], to remove as much as possible the effects of any bias we might encounter when comparing performance between FISes.

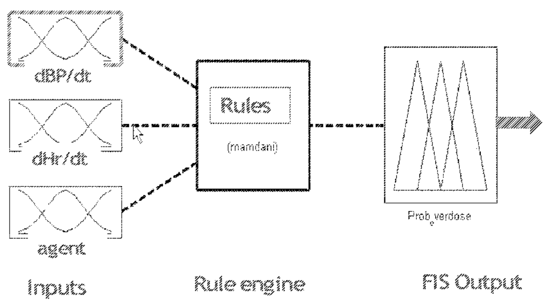

Figure 1. The feedforward FIS model designed for patient status estimation in OR room. The input parameters are trends of blood pressure and heart rate, and anesthesia agent level.

C.FIS with additional sequential inputs

To utilize sequential pattern information in an FIS we can use the first method described (Section I) and add delayed input parameters to the original FIS input as shown in Figure 2. The delay factor K can be adjusted in the model. A 2 minutes delay factor is used as default. If one additional input is used with three states then the maximum number of possible rules in the entire system increases by three-fold.

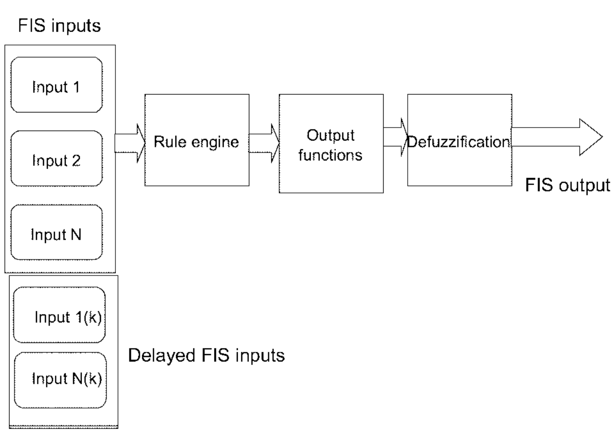

Figure 2. The feedforward FIS model designed for patient status estimation in OR room. The input parameters are trends of blood pressure and heart rate, and anesthesia agent level.

D.Recursive FIS

Using the second method for utilizing sequential pattern information we add the FIS output as an input parameter. An RFIS for sequential pattern recognition designed for estimating hemodynamic stability under inhaled anesthesia combines current status inputs including blood pressure trend, heart rate trend, anesthesia agent level with the previous output of the system is shown in Figure 3. In the system, the previous FIS output is used as feedback to the input layer with the other input parameters. The rules in the system combine feedback input and other input parameters. The feedback transfer function H(z) controls the number of previous FIS outputs, and the length of the delay taps utilized in the loop.

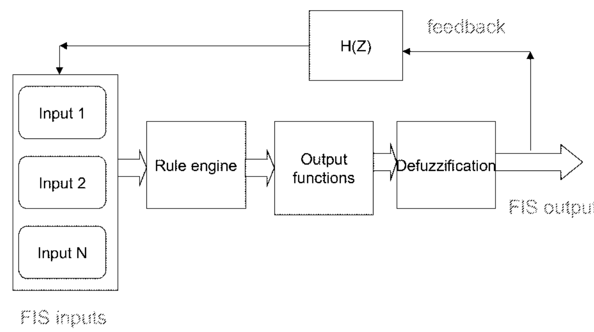

Figure 3. The RFIS model for patient status estimation in an OR room. The previous FIS output is used as a feedback to combine with the input parameters (trends of blood pressure and heart rate, and anesthesia agent level) in the rule engine.

A general structure of the multi-tap feedback transfer function of the RFIS is shown in Figure 4, where one of the rule sections is displayed. In the selected rule section, multiple feedbacks at different taps are combined with input parameters. In the structure shown, $\{b_j\}$represent membership functions of the input, and $\{a_i\}$represent membership functions of the feedback; F{ } represents the rule set of the system.

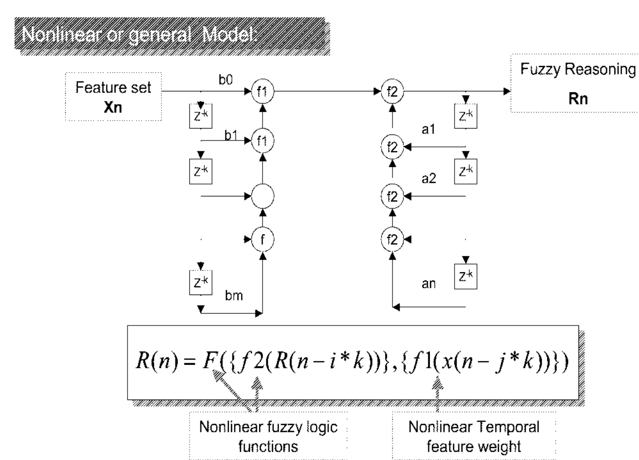

Figure 4. One of the combination branches in the RFIS model, b_js represent membership functions of the input, and a_is represent membership functions of the feedback; F{ } represents the rule engine of the system.

The RFIS needs to be set to an initial state. The form of the feedback H(Z) can be as simple as a one tap delay z^{-1}, an m step delay z^{-m}, or a sequence of delays. In addition to direct feedback, we can transform the FIS output before using it as feedback. For example, the mean or median of previous M FIS outputs. A weight or a group of weights can be added to feedback FIS outputs to control the sensitivity of the feedback outputs. For temporal pattern recognition, the goal for an RFIS system is to reinforce the similar sequential patterns to reduce the false alarm rate and at the same time to improve the sensitivity.

To create an RFIS for evaluation, we started with the same three non-recursive inputs used by the baseline FIS and added a recursive input. The rule set now has four inputs, the original three, plus the FIS output, and 27 possible rules in its rulebase. We selected a delay of 120 seconds to the previous output feedback. The value of 120 seconds was selected qualitatively, after several trial runs with different applied delays. Otherwise, the same procedures for creating and evaluating the baseline FIS were applied.

E. FIS Evaluation

To characterize the performance of each FIS we used a subset of the collected data (Section II A) and executed the FIS against those data sets and collected the FIS output. The test data set contains over 81 hours of data with almost 4 hours of data annotated as anesthetic overdose. The high number of hours annotated as anesthetic overdose is probably due to the fact the cases chosen were all high acuity surgeries. A threshold was applied to the FIS output values, values above the threshold are considered to be a positive detection of a specific condition. By comparing the detection values against the expected detections based on clinical annotations we were able to generate true positive, false positive, true negative, and false negative counts for a given threshold. By varying the threshold we can create a ROC (receiver operator curve) for the sensitivity and specificity for each threshold setting. The area under the ROC curve is used to characterize the algorithms performance.

III. Results

An example of anesthetic overdose is shown in Figure 5. It is helpful to understand the difference between the recursive and non recursive systems by looking at specific examples of the system behavior during anesthetic overdose events. In Figure 5, 2 events and the FIS output from the recursive and 4 input non recursive inference systems are shown. The top trace is the slope of mean arterial pressure which shows the falling trend (negative values) starting from 1.9 ksec, ending around 1.94 ksec into the case; the agent level is relatively high through out the period; the RFIS output (dashed line) reached 0.6 a little after the initial starting point, and increased to 0.7 as more and more feedback reinforced the possibility of anesthetic overdose. For the non-recursive output, the value peaked at around 0.6. When compared with non-recursive model output (dotted line) curve, the recursive model has a higher value in the event region due to a reinforced feedback from previous output. The FIS output remained high throughout the blood pressure falling period. The feedback from the previous output also results in the recursive FIS output rising later at the leading edge and falling slower at the trailing edge of the event than the non-recursive output. This is a result of the fact that the inference system suppresses the output in the case where there is no prior history of the condition and enhances it when the prior history indicates a high likelihood of anesthetic overdose. This effect leads to a more accurate determination of the start and end times for this particular event.

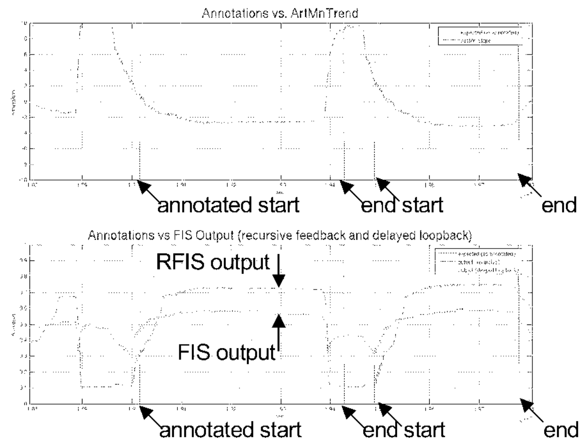

Figure 5. The case study of RFIS. The top trace is the trend of mean arterial blood pressure, and the bottom traces are RFIS output (dashed, darker) and feedforward FIS augmented with delayed inputs (dotted, lighter). The solid lines are the beginnings and the endings of agent overdose events annotated by doctors.

The test results of the RFIS and non-recursive FIS showed similar performance measured with the area under the ROC curve of approximately 0.78.

III. Discussion

This preliminary performance test used the same threshold for both FIS models, the sensitivity and specificity measure has not quite reflected the higher FIS output of RFIS model. In a further study, a higher threshold will be used to test the RFIS model to take advantage of its higher output value when sequential feedback reinforces the event. We expect the specificity could be improved when threshold is set higher to reduce false positive alarms during monitoring. In that study we also plan to use additional performance measurements to understand the nature of sequential pattern recognition when using RFIS. Such measurements might include the consistency of FIS output within a defined window and time shift of the output with respect to the annotated event. For example, 5 minutes before the annotated starting point, or 5 minutes after annotated ending point might still be considered as correct output.

If we consider a system where every input is associated with three linguistic values (for example: high, medium, low), then for each additional input that we do not need we reduce the maximum number of rules required by three-fold. This reduction in rule size has the obvious benefit that computation costs are reduced. More importantly, the fewer inputs there are the easier it is to audit the rules and compare them to existing clinical knowledge.

Current performance measurements (sensitivity vs. specificity) are based on annotated time lines of the events by clinicians retrospectively. The choice of a retrospective annotation strategy is largely motivated by constraints associated with trying to annotate in situ in the OR environment [2], but runs the risk that the annotators may be influenced by acasual information. The FIS executes casually, in such a way that only prior data information before current time is used. Measuring performance of the algorithm against the real-time performance of a clinician in the operating room might be a more accurate comparison than a comparison to retrospectively annotated data. However, the current comparisons should be sufficient to provide feedback that can be used to improve algorithm performance.

We have empirically tried this approach with several other systems and have shown tentatively that it improves performance, but depending on the problem, there may be other more immediate features which improve performance without including some form of temporal information.

As for any feedback system, we also need to consider the stability of the system. Currently the stability is established through the rules, in which any single rule output is limited to a normalized threshold.

Acknowledgment

University of Michigan Medical Center Drs O'Reilly and Rosenberg and Amy Shanks for clinical support and data collection.

References

[1]J.-S. R. Jang, C.-T. Sun, E. Mizutani, *Neuro-Fuzzy and Soft Computing*, Upper Saddle Hall, NJ 07458, Prentice Hall.

[2]M. Krajnak, J. Xue, "Optimizing Fuzzy Clinical Decision Support Rules Using Genetic Algorithms", submitted, IEEE EMBS 2006.

[3]T. Monk, V. Saini, B. Weldon, J. Sigl, "Anesthetic Management and One Year Mortality after Noncardiac Surgery", Anesthesia and Analgesia, vol. 100, 2005, pp. 4-10.

[4]J. Pendergast, et. al., "Anesthetic Overdose Input Parameter Prioritization", unpublished.

[5]D. Reich, et. al., "Predictors of hypotension after induction of general anesthesia", Anesthesia and Analgesia, vol. 101(3), 2005, pp. 622-8.

[6]B. D. Ripley, Pattern Recognition and Neural Networks, Cambridge, 1996

[7]Q. Xue, B. Taha, S. Reddy, T. Aufderheide, "An Adaptive Fuzzy Model for ECG Interpretation", presented at Marquette ECG Conference, 1998.