QUALITATIVE REASONING AND DECISION TECHNOLOGIES
N. Piera Carreté and M.G. Singh (Eds.)

CONTROLLING AN ANAESTHETIC AGENT BY MEANS OF FUZZY INDUCTIVE REASONING

Angela Nebot
Institut de Cibernètica,
Univ. Politècnica de Catalunya,
Diagonal 647,
08028 Barcelona,
SPAIN

François E. Cellier
Dep. of Electrical and Comp. Engng.,
The University of Arizona,
Tucson ,
Arizona 85721,
U.S.A.

Derek A. Linkens
Dep. of Automatic Control and Systems Engng.,
Univ. of Sheffield, P.O. Box 600, Sheffield S1 4DU,
ENGLAND

Abstract

The control of depth of anaesthesia is a difficult undertaking. Progress has been made during past years by use of different methodologies and monitoring systems that suggest the safe amount of anaesthetic drug, considering the condition of an individual patient. Despite these improvements, anaesthetists still rely heavily on personal experience (gut feeling) when suggesting the anaesthetic dosage during surgical operations. The purpose of this paper is to describe an anaesthetic agent control system using a new qualitative modeling and simulation methodology called *fuzzy inductive reasoning*.

The paper also makes a comparision between the Neural Network approach and the Fuzzy Inductive Reasoning technique when applied to biomedical applications. To this end, an Artificial Neural Network for Anaesthetic Dose determination, ANNAD [10] previously developed at the University of Sheffield, is taken as a starting point. As an alternative, a Fuzzy Inductive Reasoning model for Anaesthetic Dose, FIRAD, has recently been developed at the Technical University of Catalonia, which is discussed in this paper. The same data that were used to drive ANNAD have also been used with FIRAD in order to be able to compare the results from the two approaches.

It will be shown that the Fuzzy Inductive Reasoning technique is a valid alternative to Neural Networks also in an environment where data quality is low, such as in biomedicine. Even further, Fuzzy Inductive Reasoning guarantees that the model will not forecast behavior beyond a time for which the available data are insufficient to substantiate the prediction.

A discussion of the obtained results and an analysis of the comparative performance of the two techniques concludes the paper.

1 INTRODUCTION

Both sleep and general anaesthesia are states of unresponsiveness which vary in depth. While sleep is healthy, natural, and repeats itself rhythmically once every 24 hours, anaesthesia is an artificial state maintained by the continuing presence of chemical agents in the brain.

Anaesthetic agents affect the respiratory system, the cardiovascular system, the central nervous system, and the muscles. The use of anaesthetic agents can produce severe complications and side effects, which, under extreme conditions, may even cause the death of the patient. It is therefore essential that the dose of anaesthetic agents is limited to the

minimum amount necessary for proper anaesthesia thereby reducing undesired side effects and minimizing the risk to the patient.

Monitoring devices can be used to record the values of indicator variables, to reason about the consistency of these values, and suggest to the anaesthetist an appropriate dose of anaesthetic agent. Research results have recently been reported in the area of monitor integration that enhance the clinical robustness of such monitoring devices by improving their reasoning capabilities through the detection of critical events and by means of enhancing their alarm accuracy [11].

Several new results have been reported in the past few years relating to the control of the depth of anaesthesia. Both open–loop and closed–loop techniques have been explored ([5], [8], [12]).

One of these studies resulted in the development of a computer–based on–line expert system called RESAC (Real–time Expert System for Advice and Control) [9]. RESAC comprises a rule–based backward chaining inference engine with about 400 rules and makes use of fuzzy logic and Bayesian reasoning. The rule–base was obtained through knowledge acquisition in consultation with expert anaesthetists [5]. The major problem of this approach is the formidable size of the resulting rule–base. Obviously, this has to be a real–time expert system in order to be of any practical use.

Triggered by the aforementioned difficulties, another study was carried out by the same group that promised to enhance the run–time efficiency of the monitoring system. The new system, ANNAD, involves an Artificial Neural Network for Anaesthetic Dose determination [10]. ANNAD is a feedforward neural network trained through back–propagation. This work is reviewed in the next section.

2 BACKGROUND: ANNAD

The artificial neural network approach was chosen due to its inherent ability to learn the input/output behavior of a system in situations where it is possible to specify the inputs and outputs, but where it is difficult to define analytically a relationship between them. This is precisely the situation in biomedical applications, such as anaesthesia, since clinical signals are readily available through measurements, but no precise analytical relationships are known between them, and variations between patients are large. Also, neural networks are inherently parallel in nature, and are therefore well suited for real–time environments.

The clinical variables comprising heart rate (HR), respiration rate (RR), systolic arterial pressure (SAP), gender, age, and weight of the patient were selected as the key clinical indicator signals to be used for suggesting an anaesthetic dose (control signal).

A patient model and a controller model were independently synthesized by means of the neural network methodology. The control loop was then closed as shown in Figure 1.

Artificial Neural Network Patient Model

An Artificial Neural Network (ANN) patient model was obtained using a back–propagation algorithm applied to a set of data measured on a patient during a surgical operation.

Three separate neural networks were trained, one for each output: HR, SAP, and RR.

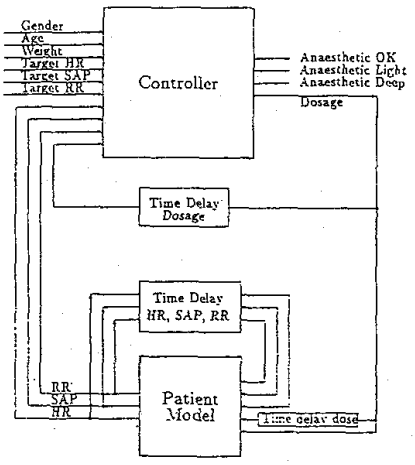

Figure 1: Feedback loop involving patient simulator and drug controller

The inputs for the training networks were the Dose, older (delayed) values of the Dose, as well as delayed values of HR, SAP, and RR. Each neural network employed two hidden layers.

Artificial Neural Network Controller

An ANN controller model (ANNAD) was obtained using a back–propagation algorithm applied to another set of data collected from a second patient during a similar surgical operation as for the patient moel. In this case, a neural network with three hidden layers was found to be optimal (showing the smallest deviation from the measured data after training). The inputs for this neural network were Gender, Age, Weight, RR, SAP, HR and the desired values of the latter three variables, while the output was the anaesthetic agent, Dose.

The anaesthetic agent used in all these operations was *isoflurane*, because the elimination of this agent by the body is quicker than that of other anaesthetic agents, such as enflurane. Obviously, both the patient model and the controller model depend heavily on the choice of the anaesthetic agent in use, and the process of model identification must therefore be repeated from scratch if another agent is to be used in the process.

Closed-loop Control

As shown in Figure 1, the control loop was then closed by connecting ANNAD with the ANN patient model. The results of this experiment demonstrate the stability of the control loop. ANNAD was able to replicate satisfactorily the advice that was obtainable from RESAC. ANNAD also produced good control performance when coupled to a patient simulator. Contrary to RESAC, which was actually used during surgical operations, ANNAD has not yet undergone real-life testing.

For a deeper insight into this work, the reader is referred to [10], [12].

3 FIRAD

The motivation for the research described in this paper was to investigate how the *fuzzy inductive reasoning* methodology performed in comparison with the neural network approach

when applied to the identification of dynamic processes from the soft sciences. To this end, we first tried to develop a fuzzy inductive reasoning model for the patient, and then to find a fuzzy inductive reasoning model for the controller. The controller model is called FIRAD (Fuzzy Inductive Reasoning for Anaesthetic Dose). The insights gained during this research effort are detailed in the following subsections.

3.1 The Methodology

As is the case of neural networks, the inductive reasoning methodology has the ability to describe (model) systems that are not well understood, that is, systems for which physical laws are only partially or not known. Contrary to the neural network approach, the inductive reasoners contain information about the likelihood of any particular state transition. This is important for model validation purposes. If the accumulated likelihood of a particular state drops below a level that can be specified by the user, forecasting will come to a halt. In this manner, the technique guarantees that the model will not forecast behavior beyond a time for which the available data are insufficient to substantiate the prediction. Also, this technique is able to enumerate all possible system behaviors that are consistent with the available knowledge (data), and can assign a measure of likelihood of occurrence to each of them.

In the fuzzy inductive reasoning approach, the qualitative systems are represented (modeled) by a special class of finite state machines called *optimal masks*, and their episodical behavior is inferred (simulated) by a technique called *fuzzy forecasting*.

A mask denotes a structural relationship between different variables, and the optimal mask is the mask that maximizes the forecasting power of the inductive reasoning process. The optimal mask is selected from a set of candidate masks by an exhaustive search process. The quality (forecasting power) of each mask is evaluated, and the one with the highest quality is returned as the optimal mask.

Once the optimal mask is found, it is possible to derive a state transition matrix from the optimal mask and the available data. The state transition matrix is a finite–state machine that lists, for each input state (i.e., each combination of input values), all possible output states together with an assessment of the likelihood of their occurrences. Once the state transition matrix has been found, a qualitative simulation can be performed by applying the *forecasting* function of the inductive reasoning methodology.

Fuzzy inductive reasoning is accomplished using SAPS–II [1], software that evolved from the General System Problem Solving (GSPS) framework [6], [13]. SAPS–II is implemented as a FORTRAN–coded function library of CTRL–C.

For a deeper insight into this methodology, the reader is referred to [1], [2], [3] and [7].

3.2 SAPS Patient Model

The patient model should be determined by the qualitative relationship between its input variable, the administered Dose, and its output variables, the clinical signals of the patient that reflect his or her body reaction to the amount of agent applied (SAP, HR, and RR).

In order to determine the patient model, we worked with the data of two different

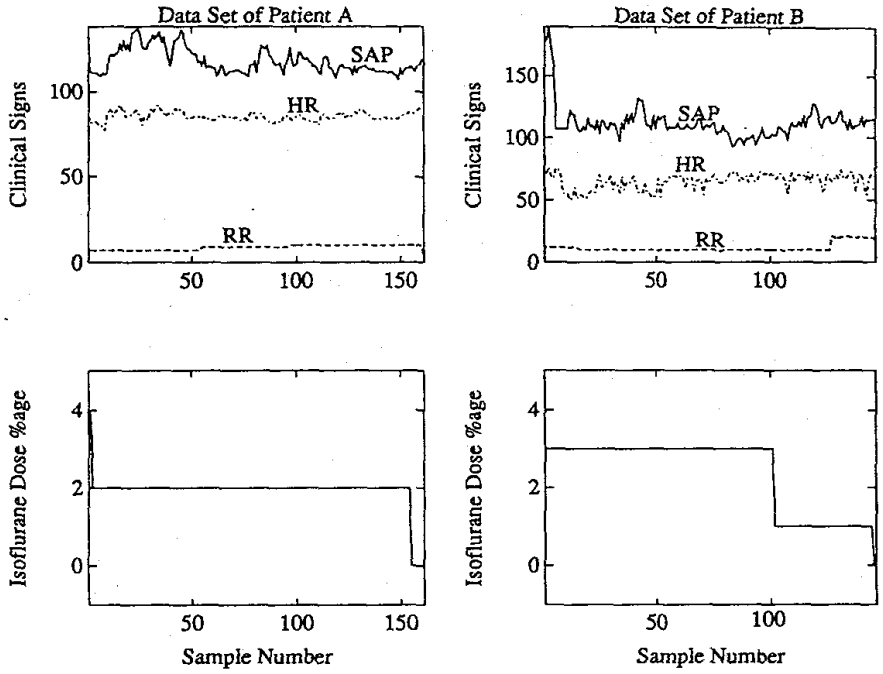

Figure 2: Patient model measurement data

patients. The available measurement data are plotted in Figure 2.

These plots reveal that the input variable, Dose, varies very little. It is "high" in the beginning of the experiment, "medium" for most part of the experiment, and goes "low" only at the very end of the experiment. It is quite clear that, in the meantime, the output variables react in various ways that are obviously not driven by the input directly, since the input does not change at all. The changes in the output variables were caused by other extraneous factors that were not recorded, and therefore, the variations in the output variables look like *noise* to the inductive reasoner. In fact, the recorded data do indeed contain considerable *digitization noise*, since all variables were recorded as integers only.

For this reason, the fuzzy inductive reasoning methodology could not find a good mask that models the patient system. The best mask synthetized did not forecast correctly. This was not the case when using the neural network methodology. It turns out that, at least for one of the data files, the neural network gave reasonable responses for the patient model.

As for all inductive techniques, inductive reasoners need a lot of data to work with. It is not possible to generate meaningful and reliable inductive models without ample and rich data. This is equally true for the neural network approach (another inductive modeling technique). However, while the neural network will always predict something, the inductive reasoner will not predict anything that cannot be validated on the basis of the available data. SAPS, our inductive reasoner, simply declines to predict anything when confronted with the patient model data, since no prediction can truly be justified given the available facts.

Here, we observe one of the *strengths* of the fuzzy inductive reasoning methodology. It will not generate models that are not justifiable from the given data. The neural network methodology generates models for *any* data, irrespective of whether they are justifiable or not. While SAPS contains an inherent model validation mechanism inside the modeling methodology, the neural network approach does not. The fact that the neural network was

slightly modified Dose of the anaesthetic drug. Therefore, there now exists a *direct causal relationship* between the observed biological data and the recommended Dose, and it should, therefore, be possible to correlate the administered Dose with the biological variables, and come up with a *causal inductive model* that can be used to replace the anaesthetist (or RESAC) in his or her (its) decision making process.

The two data sets contain 163 and 185 records, respectively. They were sampled once per minute. According to information obtained from two anaesthetists whom we consulted, the slowest time constant of interest in our system is in the order of 10 minutes, and the fastest time constant of importance is in order of 1 minute.

In accordance with Shannon's sampling theorem we should therefore use a sampling rate of approximately one sample every 0.5 minutes. For this reason, before starting to identify an optimal model, we decided to use a Spline interpolation to find one new data record per interval, located exactly in the middle between the two neighboring measurement data records. Thereby, the length of the data records was enhanced to 325 and 369 records, respectively. We can assume that not much information was lost by sampling once instead of twice per minute, because the plots reveal that the data vary sufficiently slowly.

It was then decided to recode (discretize) the variables SAP, HR, and Dose into three qualitative levels (classes), whereas RR was recoded into two qualitative levels only.

Due to the difference between the slowest and the largest time constants of importance, we decided to use a mask candidate matrix of depth 21 with nine zero rows in between rows that contain potential inputs.

$$
\begin{array}{l|cccc}
t\backslash^{*} & SAP & HR & RR & DOSE \\
t-20\delta t & -1 & -1 & -1 & -1 \\
t-19\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-11\delta t & 0 & 0 & 0 & 0 \\
t-10\delta t & -1 & -1 & -1 & -1 \\
t-9\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-\delta t & 0 & 0 & 0 & 0 \\
t & -1 & -1 & -1 & +1
\end{array}
\qquad eqno(1)
$$

In this way, one new forecast is produced every 0.5 minutes to satisfy Shannon's sampling theorem, and yet, the inductive reasoner looks at input values 5 minutes and 10 minutes back to capture the slowest time constant. This technique has proven successful in the past [4].

The first 270 (320) rows of the data matrix were used as past history data to compute the optimal mask. Fuzzy forecasting is used to predict new qualitative class and fuzzy membership values for Dose for the last 55 (49) rows of the raw data matrix, respectively.

For the first data set, the optimal mask obtained was the following:

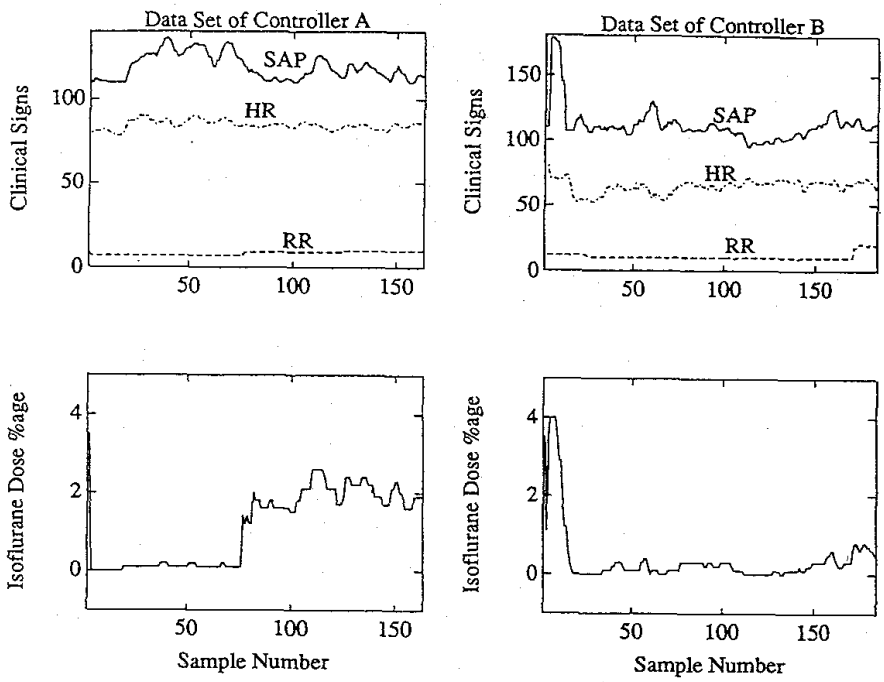

Figure 3: Controller model measurement data

able to produce a reasonable response for one of the data sets does not mean that the model is validated. The fact that it was unable to produce a reasonable response for the other data set proves just the opposite. Since inductive models necessarily lack physical insight, we believe it to be absolutely *essential* for any inductive modeler to contain an intrinsic model validation mechanism that is inseparable from the modeling tool itself. Our fuzzy inductive reasoner, SAPS, does precisely that.

3.3 SAPS controller model

The controller model is determined by the qualitative relationship between its three input variables: SAP, HR, and RR; and its single output variable: the administered Dose.

For the controller model, we were able also to obtain two data sets from two different patients. These data are plotted in Figure 3.

Looking at the plots, we can see that the output variable, Dose, varies here considerably more than in the data sets for the patient model. The patient model data are purely *clinical data*, i.e., data measured in the operating theater. The human anaesthetist didn't find the variations in the biological variables (HR, RR, and SAP) alarming, and therefore, reacted very little during the entire operation. In contrast, the controller model data were obtained from RESAC, i.e., the true biological variables that had been observed during surgery were fed into RESAC, which, in turn, proposed a value for Dose. RESAC had been validated by showing the proposed Dose to several anaesthetists, who concluded that RESAC's recommendations were clinically meaningful.

Fortunately for us, RESAC was more "industrious" than a human anaesthetist would ever be, and reacted to small variations in the biological variables by recommending a

$$
\begin{array}{c|cccc}
t\backslash^{\bullet} & SAP & HR & RR & DOSE \\
t-20\delta t & 0 & 0 & 0 & 0 \\
t-19\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-11\delta t & 0 & 0 & 0 & 0 \\
t-10\delta t & 0 & 0 & 0 & -1 \\
t-9\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-\delta t & 0 & 0 & 0 & 0 \\
t & -2 & 0 & 0 & +1
\end{array}
\quad eqno(2)
$$

This mask denotes the following relationship,

$$Dose(t) = \tilde{f}(Dose(t - 10\delta t), SAP(t)) \tag{3}$$

For the second data set, the optimal mask obtained was:

$$
\begin{array}{c|cccc}
t\backslash^{\bullet} & SAP & HR & RR & DOSE \\
t-20\delta t & 0 & 0 & 0 & 0 \\
t-19\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-11\delta t & 0 & 0 & 0 & 0 \\
t-10\delta t & -1 & 0 & 0 & -2 \\
t-9\delta t & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
t-\delta t & 0 & 0 & 0 & 0 \\
t & 0 & -3 & 0 & +1
\end{array}
\quad eqno(4)
$$

This masks denotes the relationship:

$$Dose(t) = \tilde{f}(SAP(t - 10\delta t), Dose(t - 10\delta t), HR(t)) \tag{5}$$

It turns out that the two masks obtained are *different.* Although RESAC used the same causal reasoning, SAPS decided that, by proposing a different causal relationship in the two cases, the quality of the forecast can be enhanced. The proposed controller is thus different for each of the two patients.

Since our gauge is the decision making process of RESAC, it makes sense to use the best possible mask, i.e., the mask that produces results that are as consistent as possible with those obtained from RESAC. However, both optimal masks will produce answers that a human anaesthetist would consider clinically plausible, and so, for practical purposes, it doesn't really matter which of them we use.

In reality, human anaesthetists use imprecise patient models (gut feeling) in their decision making. Thus, different optimal masks correspond to slightly inconsistent decision making, a fact that we are well prepared to accept since human decision making is never fully consistent. Different anaesthetists may decide differently when exposed to the same data, and the same anaesthetist may decide differently depending on e.g. the number of hours

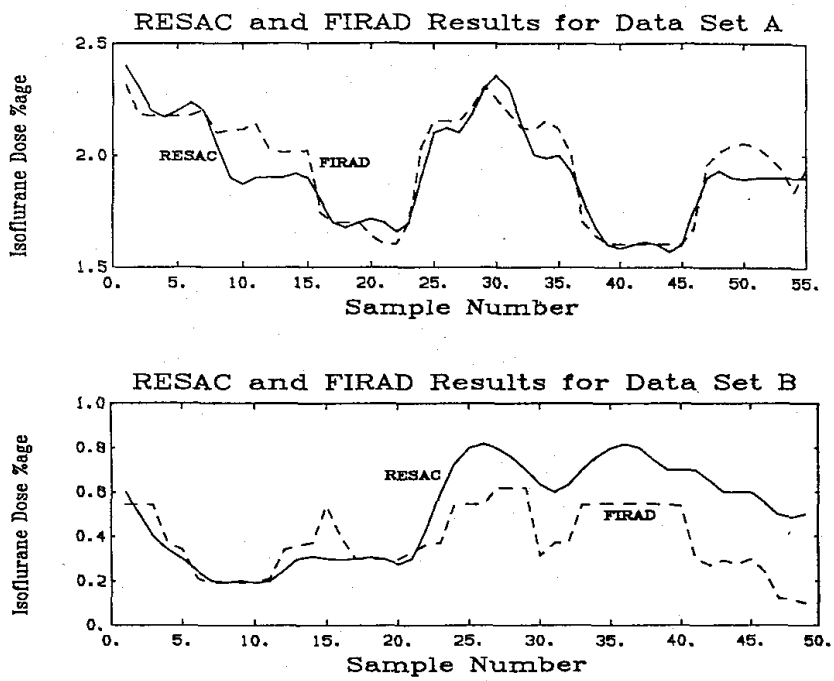

Figure 4: Comparison RESAC/FIRAD

he or she slept the night before, or depending on whether he or she had gotten into an argument with his or her spouse the evening before, or simply, depending on the current mood he or she is in. There is no such thing as a "correct" vs. an "incorrect" decision. We can only talk about "clinically acceptable" vs. "clinically unacceptable" decisions. Within the range of the clinically acceptable decisions, it doesn't matter too much how the decision is drawn. Thus, SAPS is simply being realistic in its assessment.

One fact that is common to both optimal masks is that the output of the controller model depends on the amount of previously administered anaesthetic agent. This is clinically plausible since the chemical substance accumulates in the patient for some time.

The forecast results for the two data sets are shown in Figure 4.

The results are quite good. The optimal masks contain sufficient information about the behavior of the anaesthetist (or RESAC) to be used as a valid controller of the dosage of isoflurane given to the patient. In contrast, the neural network gave good responses for the controller model only for one of the data files.

From these results, we can conclude that the SAPS methodology is fairly *robust*, i.e., it consistently generates a decent inductive model whenever the data allows it to, and it categorically will not generate a model if the available data do not permit to validate an inductive model.

The neural network approach is different in this respect, since it uses a gradient technique (backpropagation) for optimization in the original (i.e., continuous) search space, whereas SAPS uses an exhaustive search in a reduced (discrete) search space. Therefore, it is perfectly feasible that the neural network does not converge (as it happened with one of the data records), whereas SAPS will come up with the "best possible" model (within the framework of the discrete search space) whenever the data justifies a model.

Would it have been possible then to simply invert the data, i.e., use the same data records to generate a patient model as well as a controller model? The answer to this question is no. Causal modeling is an extension to the concept of uni-valued functions. Given the function $y = \sin(x)$, it is always possible to find a unique value of y for any given value of x, because $\sin(x)$ is a *uni-valued function.* On the other hand, $x = \sin^{-1}(y)$ is multi-valued, and therefore, it is not possible to conclude a unique value of x given a value of y. It therefore makes sense to call a uni-valued function a *causal function,* whereas a multi-valued function is a *non-causal function.*

Causal modeling is an extension of this concept. Causal models are uni-valued functions that accept, as inputs, not only current values of its input variables, but also past values of all its inputs *as well as its outputs.* The controller model is obviously a causal model, since the anaethetist (or RESAC) bases his or her (its) decision making in a semi-deterministic (fully-deterministic) fashion on the available inputs. The reverse, however, is not true. It is not evident that it is possible to conclude the current (and future) value(s) of the biological parameters in a unique fashion from measurements of their own past, and from current (future) as well as past Dose values, and SAPS indeed concludes that this is not a meaningful proposition.

4 COMPARISON OF RESULTS FROM THE TWO MODELING METHODOLOGIES

Before comparing the results obtained from ANNAD, and FIRAD, we wish to make a comment about the SAPS methodology. The original idea was that FIRAD should forecast the Dose during 63 minutes in order to obtain the same plot length as was shown in the previously published ANNAD report [10]. This was not possible because SAPS needed more data points (training data) than the neural network methodology to identify a model.

Previous investigations involving SAPS have led to a recommendation with respect to the minimum number of data records to be used in the identification of an inductive model. This rule is based on statistical considerations, and states that, in any class analysis, each (discrete) state should be recorded at least five times. Thus:

$$n_{rec} \geq 5 \cdot \prod_{\forall i} k_i \tag{6}$$

where n_{rec} denotes the total number of recordings, i.e., the total number of observed states, i is an index that loops over all variables, and k_i denotes the number of levels (i.e., discrete class values) of the variable i.

In the given application and using the first data stream, the number of suggested records is:

$$n_{rec} \geq 5 \cdot (3 \cdot 3 \cdot 3 \cdot 2) = 270 \tag{7}$$

Consequently, the first 270 data records should be used for model identification, which leaves us with only 55 records, or 27 minutes worth of measurement data for forecasting.

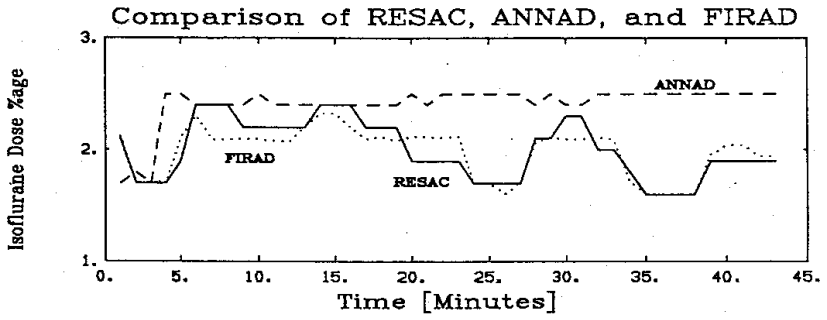

Figure 5: Comparison RESAC/ANNAD/FIRAD

To improve the situation, tests were made to find the minimum number of records needed to identify the same controller model that was found using the set of 270 records. It was determined that, if at least 240 records were used for identification, the same controller model could still be found. This then allows us to forecast the system over the last 43 minutes of the recorded data. This forecast can be compared with the forecast obtained from ANNAD and with the original Dose recommendations made by RESAC. The comparative results are given in Figure 5.

As can be seen from this plot, FIRAD forecasts the Dose quite well, even somewhat better than ANNAD. Thus, the fuzzy inductive reasoning methodology has been shown to be able to synthetize inductive biomedical models at least as well as a neural network.

5 CONCLUSIONS

The results shown in this paper confirm the ability of the Fuzzy Inductive Reasoning methodology to produce good control performance of the anaesthesic agent. The FIRAD system not only replicates the advice from RESAC, but it performs even better, more reliably, and more consistently than the ANNAD system.

As demonstrated in this paper, one of the strengths of SAPS is that it contains an inherent model validation mechanism inside the modeling methodology, which the neural network approach does not. This is why the SAPS methodology would not generate a patient model that would not have been justifiable from the given data. We consider this intrinsic model validation mechanism a *distinct advantage* in comparison with the neural network methodology, especially in the context of soft sciences.

For the controller model, two different masks have been found from two different data sets. Although RESAC used the same causal reasoning, SAPS decided that, by proposing a different causal relationship in the two cases, the quality of the forecast can be enhanced.

The SAPS methodology is fairly robust, i.e., it consistently generates a decent inductive model whenever the data permits it to do so, and it categorically will not generate an inductive model if the available data are insufficient in either quantity or quality for validation of such a model.

6 REFERENCES

1 CELLIER, F.E. and YANDELL, D.W. - "SAPS–II: A New Implementation of the Systems Approach Problem Solver,"
Int. Journal of General Systems, 13 (4), 307-322 (1987).

2 CELLIER, F.E. - *Continuous System Modeling*. Springer-Verlag, New York, 1991.

3 CELLIER, F.E. - "Qualitative Modeling and Simulation — Promise or Illusion," Proceedings 1991 Winter Simulation Conference, Phoenix, AZ, 1991, p. 1086-1090.

4 CELLIER, F.E., NEBOT, A., MUGICA, F. and DE ALBORNOZ, A. - "Combined Qualitative/Quantitative Simulation Models of Continuous-Time Processes Using Fuzzy Inductive Reasoning Techniques," Proceedings SICICA'92, IFAC Symposium on Intelligent Components and Instruments for Control Aplications, Málaga, Spain, 1992, p. 589-593.

5 GREENHOW, S.G. - "A Knowledge Base System for the Control of Depth of Anaesthesia," Ph.D. Thesis, University of Sheffield, 1990.

6 KLIR, G.J. - *Architecture of Systems Problem Solving*. Plenum Press, New York, 1985.

7 LI, D. and CELLIER, F.E. - "Fuzzy Measures in Inductive Reasoning," Proceedings 1990 Winter Simulation Conference, New Orleans, La., 1990, p. 527-538.

8 LINKENS, D.A. - "Adaptive and Intelligent Control in Anesthesia,"
IEEE Control Systems Magazine, 12 (6), 6-11 (1992).

9 LINKENS, D.A., GREENHOW, S.G. and ASBURY, A.J. - "Clinical Trails with the Anaesthetic Expert Advisor RESAC," Expert Systems in Medicine, 6th Annual Meeting, London, 1990, p. 11-18.

10 LINKENS, D.A. and REHMAN, H.U. - "Nonlinear Control for Anaesthetic Depth Using Neural Networks and Regression," Proceedings IEEE International Symposium on Intelligent Control, Glasgow, Scotland, U.K., 1992.

11 NAVABI, M.J., WATT, R.C., HAMEROFF, S.R. and MYLREA, K.C. - "Integrated Monitoring Can Detect Critical Events and Improve Alarm Accuracy,"
Journal of Clinical Engineering, July/August, 295-306 (1991).

12 REHMAN, H.U., LINKENS, D.A. and ASBURY, A.J. - "Dynamic Nonlinear Modeling and Control of Depth of Anaesthesia for Spontaneously Breathing and Ventilated Patients,"
submitted to: *Computer Methods and Techniques in Biomedicine.*

13 UYTTENHOVE, H.J. - "SAPS — System Approach Problem Solver," Ph.D. Dissertation, SUNY Binghampton, N.Y., 1979.