Rutledge G., Thomsen G., Beinlich I., Farr B., Sheiner L., Fagan L.M.: Combining Qualitative and Quantitative Computation in a Ventilator Therapy Planner, Proceedings 13th Annual Symposium on Computer Applications in Medical Care (SCAMC–89), Washington, DC, pp.315–19, 1989.

Rutledge G.W., Thomsen G.E., Farr B.R., Tovar M.A., Polaschek J.X., Beinlich I.A., Sheiner L.B., Fagan L.M.: The Design and Implementation of a Ventilator–management Advisor, Artificial Intelligence in Medicine, 5(1), pp.67–82, 1993.

Shortliffe E.H.: Knowledge–Based System in Medicine, Proceedings of the Medical Informatics Europe 1991 (MIE–91), Springer, Berlin, pp.5–9, 1991.

Uckun S., Dawant B.M. : Qualitative Modeling as a Paradigm for Diagnosis and Prediction in Critical Care Environments, Artificial Intelligence in Medicine 4(2), pp.127–44, 1992.

Uckun S., Dawant B.M., Lindstrom D.P.: Model–based Diagnosis in Intensive Care Monitoring: the YAQ Approach, Artificial Intelligence in Medicine, 5(1), pp.31–48, 1993.

Uckun S.: Intelligent Systems in Patient Monitoring and Therapy Management, Stanford University, Knowledge Systems Laboratory, Report KSL 93–32, 1993.

tive and absolute time representation (e.g. "adjustment of the ventilator setting PIP was performed 10 minutes before the neonate's health condition got better", "to improve the ventilation takes longer than to improve the oxygenation"). Additionally, we are improving our component of assessment of different kinds of trends in the sense that we will apply different trend–detection methodologies.

Acknowledgement

The current phase of the project is supported by the "Jubiläumsfonds der Österreichischen Nationalbank", Vienna, Austria, project number 4666. We also greatly appreciate the support given to the Austrian Research Institute for Artificial Intelligence (ÖFAI) by the Austrian Federal Ministry of Science and Research, Vienna.

References

Ash D., Gold G., Seiver A., Hayes–Roth B.: Guaranteeing Real–time Response with Limited Resources, Artificial Intelligence in Medicine, 5(1), pp.49–66, 1993.

Chambrin M.–C., Chopin C., Ravaux P., Mangalaboyi J., Lestavel P., Fourrier F. : RESPAID: Computer–Aided Decision Support for Respiratory Data in ICU, in Proceedings of the Eleventh Annual Conference IEEE Engineering in Medicine and Biology Soc., Seattle, WA., 1989.

Doyle R.J., Sellers S.M., Atkinson D.J.: A Focused, Context–Sensitive Approach to Monitoring, in Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI–89), Morgan Kaufmann, Los Altos, CA, pp.1231–1237, 1989.

Garfinkel D., Matsiras P.V., Lecky J.H., Aukburg S.J., Matschinsky B.B., Mavrides T.G.: PONI: An Intelligent Alarm System for Respiratory and Circulation Management in the Operating Room, in Proceedings of the 12th Annual Symposium on Computer Applications in Medical Care (SCAMC–88), Washington, DC, 1988.

Haimowitz I.J., Kohane I.S.: Automated Trend Detection with Alternate Temporal Hypotheses, in Bajcsy R.(ed.), Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI–93), Morgan Kaufmann, San Mateo, CA, pp.146–151, 1993.

Hayes–Roth B., Washington R., Ash D., Hewett M., Collinot A., Vina A., Seiver A.: Guardian: A Prototype Intelligent Agent for Intensive–Care Monitoring, Artificial Intelligence in Medicine, 4(2), pp.165–66, 1992.

Hayes–Roth B.: Intelligent Control, Special Volume: Artificial Intelligence in Perspective, Artificial Intelligence, 59(1–2), pp.213–20, 1993.

Miksch S., Horn W., Popow C., Paky F.: VIE–VENT: Knowledge–Based Monitoring and Therapy Planning of the Artificial Ventilation of Newborn Infants, in Andreassen S., et al. (eds.), Artificial Intelligence in Medicine: Proceedings the 4th Conference on Artificial Intelligence in Medicine Europe (AIME–93), Munich, pp.218–29, 1993.

Neitzke M.: Modeling Physical Systems with Relative Descriptions of Parameters, in Neumann B.(ed.), Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI–92), Wiley, Chichester, UK, pp.683–684, 1992.

For example, the degree of artificial ventilation determined by values of the ventilator settings can lead to modification of the transformation process. If the peak inspiratory pressure (PIP, measured in cm H_2O) is very high, extremely deranged $PtcCO_2$ values are tolerated as better ones.

E.g., IF (30 < PIP $\leq$ 35) and
(PtcCO$_2$ is ”extremely below target range”)
THEN (PtcCO$_2$ is changed to ”substantially below target range”)

IF (PIP > 35) and
(PtcCO$_2$ is ”extremely below target range”)
THEN (PtcCO$_2$ is changed to ”slightly below target range”).

5. Technical Remarks and Evaluation

VIE–VENT was implemented using the knowledge representation language Clips (NASA). We used forward chaining rules for representing the knowledge base. VIE–VENT is running on IBM–compatible personal computers, Apple Macintosh and UNIX–workstations.

VIE–VENT samples transcutaneous measurements and ventilator settings every 10 seconds. These data are used for data validation, data abstraction and data–oriented therapy recommendations in critical situations. In addition they provide the basis for calculating very short–term and short–term trend data. These data are used to compare observed values with expected behavior. The arithmetic means of these 10–second data are stored every 10 minutes for in-depth analysis using medium–term and long–term trends. These data provide the basis for goal–oriented therapy recommendations (especially during weaning) and to preserve a stable clinical situation where parameters stay within ”normal” limits.

So far, we did not evaluate every module separately. But we performed a technical evaluation of the whole VIE–VENT system on real problems. The knowledge engineer and two domain experts participate in the evaluation. Our sample consisted of seven real cases from neonates' case histories and six generated cases by the two physicians, each covering several prototypical and extreme situations. We observed 22320 decision steps. The physicians ranked VIE–VENT's therapy recommendations, warnings and explanations as ”correct”, ”correct, but needs smoothing” and ”incorrect”. 65% were ranked as ”correct”, 31% as ”correct, but needs smoothing” and 4% as ”incorrect”. The ”incorrect” recommendations were caused by the fact that extreme situations are underrepresented in the knowledge base and by the lack of taking into consideration long–term effects sufficiently. Up to now, evaluation showed no problems with our data validation and data abstraction components. The physicians appreciated our idea of dealing with missing values, the priority lists, the dynamic calibration of values, the context–sensitive adjustment of qualitative values and the assessment of different kinds of trends for data validation and therapy planning.

Future enhancements of VIE–VENT cover different levels. VIE–VENT represents a data–driven approach with only limited time representation. Therefore, we are currently expanding our model of artificial ventilation (the neonatal respiration) with a combination of relasite of measurement: **arterial**

KIND OF BLOOD GAS	below target range			Target Range			above target range			mode of ventilation
	s3	s2	s1		TV		g1	g2	g3	
pO_2	30	45	55	60	65	70	75	120	700	
SaO_2	20	80	87	91	93	94	97	100	100	
pCO_2	130	55	40	37	35	32	30	20	15	IPPV
pCO_2	130	70	52	45	40	37	35	30	15	IMV
pH	6,7	7,2	7,28	7,3	7,4	7,45	7,5	7,65	7,8	

Explanation of abbreviation:

below target range
s1 ... slightly
s2 ... substantially
s3 ... extremely

TV ... Target Value

above target range
g1 ... slightly
g2 ... substantially
g3 ... extremely

Table 1: Transformation schema of arterial blood gas measurements

4.2.2 Dynamic Calibration of Values

The (static) transformation schemata provide a very clear and useful method for achieving qualitative values. However, in a complex situation with several monitor parameters and even several monitors, dynamic calibration of these schemata is necessary. When we observe several sensor channels, we know about parameters which reflect the identical patient's situation. If these parameters deviate from each other due to the individual situation of the patient or due to variations in the environment conditions under which the sensor operates, we need a dynamic adjustment. This is done under the assumption that the data validation task has classified the data as reliable. The method we use is linear calibration based on the reliability ranking. The condition of activating calibration depends on the measurement. In VIE–VENT calibration is only done in case the qualitative values differ by two qualitative categories.

As an example: One of the benefits of VIE–VENT is the opportunity to combine values of the transcutaneous blood gas monitor with discontinuous blood gas measurements. However, transcutaneous blood gases may deviate from the more reliable but only rarely drawn invasive blood gas samples. Therefore transcutaneous measurements have to be calibrated against invasive blood gas measurements. Although there is no tight linear correlation between the two types of measurement, we decided to use a linear calibration factor for practical reasons (e.g., $k = (PCO_2 / PtcCO_2)$; $PtcCO_2(new) = k * PtcCO_2$ (actual)). Calculated values are transformed into the qualitative values, which are used in our system model of neonatal respiration.

4.2.3 Context–sensitive Adjustment of Qualitative Values

For extremely critical or life–threatening situations of a patient, the thresholds defined in the transformation schemata are too strict. In these cases we adjust the qualitative values of a parameter, which is equal to a shift of the numerical threshold value.

When all measurements are "unknown" or a critical situation has arisen in the past, VIE–VENT is unable to find a solution and the recommendations of appropriate treatments are shifted to the physician. In other situations VIE–VENT recommends changes or maintenance of the degree of artificial ventilation depending on the parameters of the ventilator and the "priority lists" (see chapter 4.1.3). Using default rules is a weak method, because it requires a large number of rules. For each combination of known and unknown parameters a separate rule is needed.

4.2 Data Abstraction

Data abstraction is the process of transforming quantitative data of the observable system into qualitative values. The qualitative values of the parameters are used in the system model for data interpretation and therapy planning. We distinguish between three different kinds of data abstraction: the absolute transformation of quantitative data into qualitative values, the dynamic calibration of values and the context–sensitive adjustment of qualitative values.

4.2.1 Absolute Transformation of Quantitative Data into Qualitative Values

The absolute transformation of quantitative data into qualitative values is usually performed by dividing the numerical range of a parameter into regions of interest. Each region stands for a qualitative value. This region defines the only common property of the numerical and qualitative values.

In VIE–VENT the basis of the transformation of the blood gas measurements are schemata, which result in seven qualitative categories of the degree of blood gas abnormalities. These schemata are defined for all kinds of blood gases depending on the sampling site of measurement (arterial, capillary, venous, transcutaneous) and the mode of ventilation (IPPV, IMV).

Table 1 shows one example of such a transformation schema of the arterial blood gas measurements during intermittent positive pressure ventilation (IPPV) or intermittent mandatory ventilation (IMV). The middle of the table indicates the expected normal value range, the target range. "TV" is the target value we are aiming to reach. The term "below target range" means that the amount of artificial ventilation is too low. The term "above target range" means that the amount of artificial ventilation is too high. For example, the transformation of the arterial pCO_2 value of 47 mmHg, when the mode of ventilation is IMV will result in a qualitative pCO_2 value of s1 ("slightly below target range"). An advantage of using qualitative values is their unified usability in the system model, no matter of which origin they are. Adaptation to specific situations can easily be done by using specific transformation tables without changing the model of respiration.

($E(Y) = \alpha + \beta * X_i$, where $E(Y)$ is the expected value, X_i are the observed data points, α is a constant value (offset), and β is the growth rate). We assume that the observations are mutually independent and have the same variance.

Firstly, the growth rate allows to classify a measurement as being increasing or decreasing. This supports the use the functional dependency checks on very short–term and short–term trends.

Secondly, for each measurement there are limits which define an implausible high increase or decrease respectively. If the growth rate β exceeds this limit a sensor problem is assumed, the data value is invalidated and an alarm is triggered.

Thirdly, we compare changes of directions of the growth rate of the very short–term and the short–term trends. Let β_v be the growth rate of the very short–term trend and β_s be the growth rate of the short–term trend. If $sign(\beta_v) = sign(\beta_s)$, with $|\beta_s|, |\beta_v| > \epsilon$, then the change of direction is judged as true. This is an important information, which is forwarded to the data interpretation module. It allows to recognize if there is either a serious problem coming up or if a therapeutic action leads to the expected result.

Fourthly, an additional very relevant figure is the difference between the amount of the growth rates $d = |\beta_s - \beta_v|$. Based on two thresholds β' and β'' we identify three situations characterized by d:

if $0 \leq d < \beta'$	*then* no essential change of the measurements is assumed;
if $\beta' \leq d < \beta''$	*then* this major change of the measurements is indicating the beginning of a serious problem;
if $\beta'' \leq d$	*then* extremely rapid changes are physiologically impossible, therefore the last data values of this measurement must be artifacts.

Currently we are experimenting to find appropriate β' and β'' values for each measurement.

4.1.5 Missing Values

A robust system has to deal with missing values. In principle, there are two ways how missing values are "produced": either a value is marked as an artifact by the data validation procedure, or we receive no data from the monitor. There are two options to deal with missing values:

(a) using a simplified system model for data interpretation:
The simplified system model uses few parameters in its model. VIE–VENT uses a simplified system model of neonatal respiration during the initial phase when the only reliable continuous measurement is SaO_2. There are restricted reactions to decrease oxygenation depending on the degree of abnormality of the SaO_2 and the actual tidal volume (VT). The VT is estimated here by the extent of chest wall expansion.

(b) context–dependent rules applying defaults for missing values:
When the qualitative value of a measurement is transformed to "unknown", VIE–VENT triggers a default rule. The default rules are measurement dependent and context sensitive.

Causal and functional dependencies of different measurements are comparable to the relative descriptions of parameters in physical systems (Neitzke, 1992) in the sense that a parameter is described relative to another one.

4.1.3 Priority Lists for Reliability

Priority lists of the measurements are an indicator of the reliability of measurements. The data validation process allows to identify a less reliable parameter from a set of conflicting parameters. The result is a *reliability ranking*. From the medical and technical sampling point of view, there is a well–defined priority which measurement is more reliable than another, depending on different conditions. On the one hand these lists facilitate the data validation task and on the other hand they also help the pruning of different and concurrent therapy recommendations.

Examples of priority lists of VIE–VENT are:

arterial blood gases are more reliable than venous blood gases;
invasive blood gases are more reliable than transcutaneous blood gases;
oxygenation: PO_2 is more reliable than SaO_2 and SaO_2 is more reliable than $PtcO_2$.

4.1.4 Assessment of Different Kinds of Trends

Comparison of different kinds of trends is an appropriate method to detect rapid oscillations or very fast changes of a single measurement as well as to eliminate artifacts. We distinguish four kinds of trends based on our samples, which are derived from new measurements every 10 seconds. The distinction of the trends are guided by physiological criteria:

(a) *very short–term* trend: sample of data points based on the *last* minute (six data points)
(b) *short–term* trend: sample of data points based on the *last 10* minutes (60 data points)
(c) *medium–term* trend: sample of data points based on the *last 30* minutes (max. 62 data points)
(d) *long–term* trend: sample of data points based on the *last 3* hours (max. 77 data points)

The arithmetic means of the 10–second data are stored after every 10 minutes. Therefore a maximum of 62 data points are available for the medium–term trend (2 arithmetic means and maximum 60 of 10–second data) and a maximum of 77 data points for the long–term trends (17 arithmetic means and maximum 60 of 10–second data). The medium– and long–term trends are mainly used as basis for goal–oriented therapy recommendations.

The problem in the field of artificial ventilation of newborn infants – as in other medical fields, like pediatric growth (Haimowitz, et al. 1993) – is the lack of an appropriate curve–fitting model of the growth of measurements which could be matched with the actual measurements. Therefore our first effort is to approximate the growth of the continuously assessed measurements $PtcO_2$, $PtcCO_2$ and SaO_2 with a simple linear regression model

VENT uses the following methods for detecting artifacts: checking the plausibility of measurements, causal and functional dependencies, priority lists, and assessment of different kinds of trends.

The different methods to deal with these above–mentioned problems are discussed in the following:

4.1.1 Plausible Measurements

The most basic method is range checking like built–in hardware alarms of monitors do. We have enhanced this method by adding additional attributes, which define the clinical context (e.g. arterial, IPPV). There are look–up tables for all input parameters, which cover the plausible measurements. A parameter in the look–up table is specified by a parameter name, a list of attribute descriptors, a upper limit and a lower limit. For example, (pCO_2, (arterial, IPPV), 15, 130), where "arterial" refers to the kind of blood gas analysis and IPPV to the mode of ventilation. When a new parameter value is received, the system checks if this value is in or out of the range and a corresponding flag is set, e.g., if $15 \leq$ new_pCO_2 (arterial, IPPV) ≤ 130 then it is a plausible measurement.

4.1.2 Causal and Functional Dependencies

The causal and functional dependencies are very useful methods to detect artifacts or abnormal behavior of parameters.

A causal dependency specifies a relationship between an actual parameter and the expected value of a corresponding parameter. It is a kind of cross–relation between different measurements.

The following example shows corresponding values of

chest wall expansion		tidal volume (VT)
1 ... small	$\leftrightarrow$	$VT \leq 5$ ml/kg
2 ... normal	$\leftrightarrow$	5 ml/kg $< VT \leq 10$ ml/kg
3 ... excessive	$\leftrightarrow$	$VT > 10$ ml/kg

A functional dependency describes a functional relationship between two or more parameters. Firstly, a functional dependency may provide a value for a dependent parameter (e.g., $AMV = VT * f$, where AMV is the minute ventilation, VT is the tidal volume and f is frequency). Secondly, we use functional dependencies for checking inadequate data transmission (e.g., $f = 60 / (t_I + t_E)$, where t_I is the inspiration time and t_E is the expiration time. We receive all values of f, t_I, and t_E and check the functional dependency of these parameters). Thirdly, we can define functional dependencies for expectations on trends. The increase/decrease of one parameter suggests an increase/decrease of another one. If such an expectation is violated, one of these parameters must be faulty. E.g., if the minute ventilation (AMV) is increasing then PCO_2 is expected to decrease. Fourthly, the functional dependencies could also be used to cope with missing values. Implicitly VIE–VENT integrates this feature in the simplified system model for data interpretation (see chapter 4.1.5).

conditions (e.g., critical ventilatory condition of the neonate, elapsed time intervals). VIE–VENT uses the following input parameters:

(a) continuous data:
ventilator settings: F_iO_2, f, PIP, PEEP, t_I, t_E, vi, ve, VT
mode of ventilation: IPPV, IMV, CPAP
transcutaneous blood gases: $PtcO_2$, $PtcCO_2$, SaO_2

(b) discontinuous data:
neonate's personal description (e.g., name, sex)
clinical parameters (e.g., weight, age, chest wall expansion, spontaneous breathing effort)
invasive blood gases: pH, PO_2 , PCO_2
site of blood gas measurements: arterial, capillary, venous.

The output parameters are primarily therapy recommendations. A therapy recommendation consists of the amount and frequency of the ventilator's parameters to be changed. Additionally, VIE–VENT prints warnings in critical situations, as well as comments and explanations about the health condition of the neonate. VIE–VENT is an open–loop system.

4. Context–Sensitive Data Validation and Data Abstraction

Among other approaches, like statistical analysis and control theory, knowledge–based system technology may appropriately represent and organize the practical and theoretical knowledge of experienced specialists and help to cope with information overload with continuous data selection, data validation and therapy planning (Shortliffe, 1991). But to detect also artifacts and complex faulted behavior an adjustment of thresholds and a transformation of data into qualitative values context–sensitively and dynamically as well as a combination of statistical analysis with knowledge–based system technology is needed.

"Context–sensitivity" means validation of parameters based on the interaction of different parameters of one monitor or of several monitors and the clinical situation of the patient. In contrast, built–in hardware alarms of a monitor are simple range–checks of one parameter.

In the next sections we go more deeply into our concepts of data validation and data abstraction.

4.1 Data Validation

The major aim of the data validation process is to arrive at reliable measurements. There are different kinds of data validation actions: the checking of plausible measurements, the handling of missing values, and the process of recognizing artifacts. There are well–known methods to deal with the first two issues. But the recognition of artifacts is a rather complicated task. An artifact is a situation where the measured values do not reflect the clinical context. Several monitors have a built–in module for recognizing unusual data values, especially those arising from hardware problems. But these built–in modules often trigger a false alarm. VIE–VENT recognizes such alarms reported by the monitors, in addition. VIE–

During the past decade, several knowledge–based systems were introduced to support clinicians with the monitoring of critical care patients and to assist them in diagnostic decisions and therapy planning (Uckun, 1993). These systems range from simple intelligent alarms (e.g., RESPAID (Chambrin, at al. 1989), PONI (Garfinkel, et al. 1988)) to sophisticated systems for anesthesia monitoring or ventilator management (e.g., VentPlan (Rutledge, et al. 1989; 1993), SIMON (Uckun, et al. 1992; 1993), GUARDIAN (Hayes–Roth, et al. 1992; Ash, et al. 1993)).

Simple alarming tasks can be managed by well–known techniques, like time–series analysis or control theory. For detection of artifacts and complex faulty behavior a combination of these methods with knowledge–based system technology is needed.

3. VIE–VENT's System Architecture

Our aim in developing VIE–VENT was to incorporate monitoring and therapy planning tasks. As shown in Figure 1, the architecture of VIE–VENT consists of several modules: data selection, data validation, data abstraction, data interpretation and therapy recommendations. All these steps are involved in a single cycle of data collection from monitors. According to our design criteria of a practically oriented knowledge–based system, we built the various module components in analogy to the clinical reasoning process. VIE–VENT represents a data–driven approach. In this paper we emphasize only two components: the data validation and the data abstraction. A description of VIE–VENT's data interpretation and therapy planning task is given in Miksch, et al. (1993).

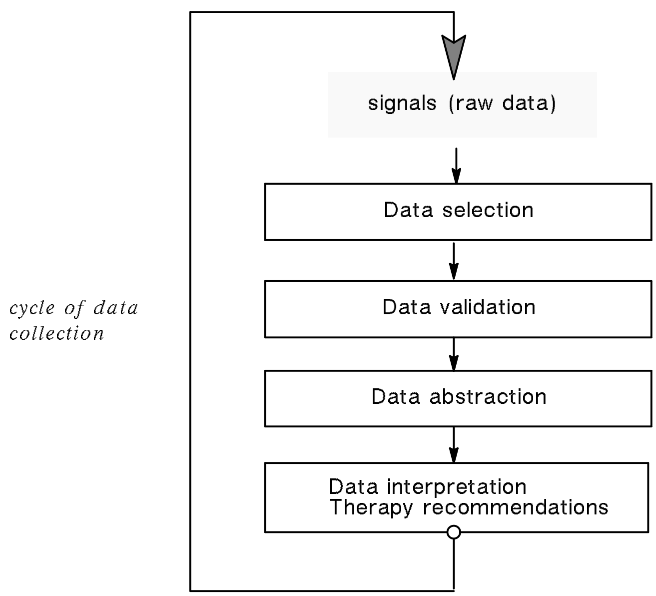

Figure 1: VIE–VENT's system architecture according to the tasks of knowledge–based monitoring and therapy planning

VIE–VENT's whole input data set can be divided into continuous and discontinuous data. The continuous data are taken from the output of the data selection module every 10 seconds. The discontinuous data are entered on request to the system depending on different account that not all sensor data can be checked in limited time and there are more faulty data than may be expected.

The monitoring problem becomes more difficult when the behavior of a system involves interactions among components or interactions with people or with the environment. Under these conditions, correct decisions become context–dependent. It is possible to determine a priori a set of sensor parameters with their fixed plausible ranges. But if the context is shifting, like one component gets in a critical condition, a capability for dynamic adjustment of threshold values is needed.

Control theory or statistical analysis is a useful task that permits a straightforward mapping between sensor data values and appropriate control action. By contrast, monitoring requires a more "intelligent" approach including tasks like interpretation and prediction of the system behavior, focussing, context–sensitive data validation and abstraction (Hayes–Roth 1993).

In this paper, we concentrate on the initial steps in the monitoring and therapy planning process – detecting anomalous system behavior quickly and arriving at reliable measurements for the therapy planning steps. Our approach is a context–sensitive focussing on relevant continuous and discontinuous data, a validating and an abstracting of these data. An important issue is the adjustment of thresholds and the transformation of data into qualitative values context–sensitively and dynamically. These components are realized in two essential modules of VIE–VENT (Miksch et al. 1993), a knowledge–based monitoring and therapy planning system of the artificial ventilation of newborn infants, which we are developing at the Austrian Research Institute for Artificial Intelligence (ÖFAI) in cooperation with the Department of Pediatrics of the Hospital of Mödling, the Neonatal Intensive Care Unit (NICU) of the Department of Pediatrics of Vienna's University Medical School, and the Department of Medical Cybernetics and Artificial Intelligence, University of Vienna.

2. The Monitoring Problem in Intensive Care Units (ICUs)

The care of critically ill patients in ICUs is increasingly complex, involving interpretation of many variables, comparative evaluation of many therapy options, and control of many patient–management parameters. The increasing demand for information storage, retrieval and processing creates problems of information management due to the increased sophistication of laboratory and monitoring equipment. The technical improvement of the ICUs' equipment makes a huge amount of measurements available to the medical staff, and even skilled physicians frequently suffer from information overload. An additional huge problem at modern ICUs is the high evidence of false alarms of the monitors during critical situations of the monitoring process.

The quality of intensive care is not only limited by the amount of information to be processed, but also by human factors, like the problem of vigilance, varying expertise, and human errors. These frequently lead to errors in diagnosis and selection of appropriate treatments.

Context–sensitive Data Validation and Data Abstraction for Knowledge–Based Monitoring

Silvia Miksch[1], Werner Horn[1 2], Christian Popow[3], Franz Paky[4]

[1] Austrian Research Institute for Artificial Intelligence (ÖFAI)
Schottengasse 3, A–1010 Vienna, Austria
email: silvia@ai.univie.ac.at

[2] Department of Medical Cybernetics and Artificial Intelligence, University of Vienna
Freyung 6/2, A–1010 Vienna, Austria

[3] Department of Pediatrics, NICU, University of Vienna
Währinger Gürtel 18–20, A–1090 Vienna, Austria

[4] Department of Pediatrics, Hospital of Mödling
Weyprechtgasse 12, A–2340 Mödling, Austria

Abstract

This paper addresses two important components of our knowledge–based system, VIE–VENT, a monitoring and therapy planning system for artificially ventilated newborn infants: data validation and data abstraction. VIE–VENT is specifically designed for practical use under real–time constraints in Neonatal Intensive Care Units (NICUs). Monitoring includes observing and guiding the behavior of a system. We concentrate on the initial steps in the monitoring and therapy planning processes – detecting anomalous system behavior quickly and arriving at reliable measurements for the therapy planning steps. The monitoring task must take into account that not all sensor data can be checked in available time and there are more faulty data than may be expected.

Our approach is a context–sensitive focussing on relevant continuous and discontinuous data, a validating and an abstracting of these data. Important issues are the detection of artifacts, the adjustment of thresholds and the transformation of numerical data into qualitative values. These methods were applied in a context–sensitive and dynamic way by recognizing the interaction between different measurements in the context of the current clinical situation of the neonate. Additionally, we used a combination of statistical analysis with knowledge–based system technology.

This paper presents a summary of the methods used for data validation and data abstraction. The methods are illustrated by examples from our VIE–VENT application.

1. Introduction: The Monitoring Problem

Monitoring involves observing and guiding the behavior of a system with real–time constraints. Monitoring consists of a number of problem–solving tasks, like recognizing abnormal conditions, combining sensor information into a picture of the global state of a system, isolating faults, predicting both normal and faulted behavior, and maintaining safe operation in the presence of faults (Doyle, et al. 1989). Additionally, decisions must often be made in limited time, and with partial information. The monitoring task must take into