

UTILIZING TEMPORAL DATA ABSTRACTION FOR DATA VALIDATION AND THERAPY PLANNING FOR ARTIFICIALLY VENTILATED NEWBORN INFANTS

Silvia Miksch¹⁾

Werner Horn²⁾

Christian Popow³⁾

Franz Paky⁴⁾

1) Austrian Research Institute for Artificial Intelligence (ÖFAI)
Schottengasse 3, A-1010 Vienna, Austria
phone: ++43 - 1 - 53 53 281 - 0
fax: ++ 43 -1 - 53 20 652
email: silvia@ai.univie.ac.at

2) Department of Medical Cybernetics and Artificial Intelligence, University of Vienna
Freyung 6, A-1010 Vienna, Austria
phone: ++43 - 1 - 53 53 281 - 0
fax: ++ 43 -1 - 53 20 652
email: werner@ai.univie.ac.at

3) NICU, Division of Neonatology, Department of Pediatrics, University of Vienna
Währinger Gürtel 18-20, A-1090 Vienna, Austria
phone: ++43 - 1 - 40 400 - 32 32
fax: ++ 43 -1 - 40 400 - 31 65
email: christian.popow@vm.akh-wien.ac.at

4) Department of Pediatrics, Hospital of Mödling
Weyprechtgasse 12, A-2340 Mödling, Austria
phone: ++43 - 2236 - 204 - 0
fax: ++ 43 - 2236 - 204 - 111

in: *Artificial Intelligence in Medicine*, 8(6), pp. 543-576, 1996.

Abstract: Medical diagnosis and therapy planning at modern intensive care units (ICUs) have been refined by the technical improvement of their equipment. However, the bulk of continuous data arising from complex monitoring systems in combination with discontinuously assessed numerical and qualitative data creates a rising information management problem at neonatal ICUs (NICUs).

We developed methods for data validation and therapy planning which incorporate knowledge about point and interval data, as well as expected qualitative trend descriptions to arrive at unified qualitative descriptions of parameters (temporal data abstraction). Our methods are based on schemata for data-point transformation and curve fitting which express the dynamics of and the reactions to different degrees of parameters' abnormalities as well as on smoothing and adjustment mechanisms to keep the qualitative descriptions stable. We show their applicability in detecting anomalous system behavior early, in recommending therapeutic actions, and in assessing the effectiveness of these actions within a certain period. We implemented our methods in VIE-VENT, an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants. The applicability and usefulness of our approach are illustrated by examples of VIE-VENT. Finally, we present our first experiences with using VIE-VENT in a real-clinical setting.

5.1 CHARACTERISTICS OF THE PROBLEM DOMAIN: REQUIREMENTS FOR A TEMPORAL-ABSTRACTION METHOD

Technical advances in the equipment used at Intensive Care Units' (ICUs') have resulted in vast increases in the volume of data available to the medical staff. Nevertheless, the care of critically ill patients is rather complex, involving the interpretation of many continuously and discontinuously assessed variables during different time periods and the risk of introducing possible fatal errors.

Artificial ventilation has greatly contributed towards the improvement of the mortality and morbidity of premature newborn infants [Goldsmith and Karotkin, 1993; Perlman et al., 1995]. Enhanced knowledge about the pathophysiological mechanisms of barotrauma and oxygen toxicity led to the development of patient-tailored strategies of mechanical ventilation and helped to reduce harmful side effects of respirator therapy. However, the bulk of continuous data arising from complex monitoring systems in combination with discontinuously assessed numerical and qualitative data creates a rising information management problem at neonatal ICUs (NICUs).

Monitoring and therapy planning with respect to real-world medical environments entails non trivial data analysis problems. The available data occur at various observation frequencies (e.g., high or low frequency data), at various regularities (e.g., continuously or discontinuously assessed data), and at various types (e.g., qualitative or quantitative data). The monitoring and therapy planning process has to cope with a combination of all these data sources. Adequate and sufficient therapy planning is based on reliable data. However, the observed data include an unexpectable high volume of faulty data (compare sample case in section 5.8). Theories on data analysis [Goebel and Schader, 1992; Kay, 1993] mostly deal with well defined problems. However, in many real-world cases the underlying structure-function models are poorly understood or not applicable because of their incomplete knowledge and complexity as well as the vague qualitative data involved (e.g., qualitative expected trend descriptions). Therefore, statistical analysis, control theory, or other techniques are often impossible, inappropriate or at least only partially applicable. Finally, the needs for quality control and quality assurance ([EURISIC, 1995], European User Requirements for Intensive Care) as well as for reusability and maintenance of knowledge [Barahona and Christensen, 1994] are increasing.

Consequently, the two most critical aspects of effective knowledge-based data analysis are data validation and therapy planning. The aim of the data validation process is to arrive at reliable measurements. Therapy planning includes three main tasks: selecting which therapeutic actions may improve the patient's condition, predicting short- and long-term outcomes of therapeutic actions, and adopting a therapeutic plan according to some explicitly defined preferences on the predicted condition of the patient [Stefanelli, 1992].

Nowadays, data validation and therapy planning do not work well. Several monitors have built-in modules for recognizing faulty data, especially those arising from hardware problems. However, these built-in modules often trigger false alarms. Monitoring data are observed by the trained medical staff. However, these single observations may only be recognized for being "normal" or "abnormal". Information about trends, "natural" oscillations, etc. is very difficult to gather. Therefore, inexperienced personnel may have difficulties in interpreting a clinical picture from pure monitoring data. Some of the variables are influenced by other clinical variables, which may not be (continuously) determined (like cardiac output, pulmonary perfusion). Complex respiratory variables, like perfusion or oxygenation, often cannot reliably be predicted from blood gas measurements. Moreover, time-series analysis techniques [Avent and Charlton, 1990] and other theories on data analysis [Goebel and Schader, 1992; Kay, 1993] may be insufficient in this domain because of the absence of an appropriate curve-fitting model and incomplete knowledge.

We present temporal abstraction methods, which are based on numerical and qualitative data with different granularity in order to arrive at unified qualitative descriptions. An advantage of qualitative descriptions is their uniform usability in the system model, irrespective of their origin. Applying these qualitative values ease data validation, data interpretation, and therapy planning. Finally, the resulting system model is easily comprehensible and transparent.

Our methods are oriented towards, but not limited to, the domain of artificial ventilation of newborn infants. They require domain specific knowledge (e.g., functional and causal dependencies, expected qualitative trend descriptions), but the principles are easily applicable in other domains with similar problem characteristics by instantiating the domain knowledge.

The paper is organized as follows. In section 5.2 we explain why other methods proposed in the literature fail to meet our specific requirements. In section 5.3 we describe the central aims and the system architecture of VIE-VENT, an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants. In section 5.4 we introduce the available knowledge. In section 5.5 we focus on the temporal data abstraction process and explain how to derive uniform qualitative descriptions based on time point and time interval data as well as on expected qualitative trend descriptions. In section 5.6 and 5.7 we show how our temporal abstraction methods can be utilized for data validation and therapy planning. In section 5.8 we present our visualization of data, gathering all available and derived information within a global and comprehensive picture. Finally, we describe our preliminary evaluation of the system within a real-clinical setting and specify the requirements that underline the applicability of our approach to other domains.

5.2 NECESSITY TO DEVELOP A NEW DATA ABSTRACTION METHOD: WHY PREVIOUS METHODS FAIL

The most common methods used in the interpretation of continuous data are time-series analysis techniques [Avent and Charlton, 1990]. Furthermore, probabilistic and fuzzy classifiers are useful approaches for classifying values. However, all of these methods have crucial shortcomings. In the absence of an appropriate curve-fitting model, time-series analysis techniques are not applicable. Domain specific characteristics, like dynamically changing degrees of parameters' abnormalities depending on the changing states of the environment cannot easily be integrated in probabilistic or fuzzy classifiers. Classifications by value interval assignments (point data abstraction methods, e.g., [Dean and McDermott, 1987]) are insufficient in dynamically changing environments where the temporal dimension covering the course of a parameter and the interdependencies of different parameters over time has to be taken into account. The shortcomings of purely mathematical methods suggest the application of knowledge-based approaches.

[Haimowitz, 1994; Haimowitz et al., 1995] have developed the concept of trend templates in their system *TrenDx* to represent all available information during a sequence of temporally discrete observation processes. A trend template defines disorders as typical patterns of relevant parameters. These patterns consist of a partially ordered set of temporal intervals with uncertain endpoints. The

trend templates are used to detect trends in series of time-stamped data. The drawbacks of this approach lie in the predefinition of the expected normal behavior of parameters during the whole observation process and the usage of absolute value thresholds matching a trend template. The absolute thresholds do not take into account the different degrees of parameters' abnormalities. In the field of artificial ventilation of newborn infants it is impossible to define such static trajectories of the observed parameters in advance. Depending on the degrees of parameters' abnormalities different expected normal behaviors are possible. These normal expectations deviate according to the patient's respiratory status in the past. Therefore, these thresholds need to be dynamically derived during the observation period. For example, the decreasing of $P_{tc}CO_2$ from 94 mmHg to 90 mmHg during the last 25 minutes would be assessed as "decrease too slow" because the patient's respiratory status was well above the target range in the past. But the same amount of change (4 units) from 54 mmHg to 50 mmHg during the last 25 minutes would be assessed as "normal decrease" if the patient's respiratory status was slightly above the target range.

Additionally, the main purpose of *TrenDx* is trend detection. It does not cover the activation of therapeutic actions and the assessment of the benefits of therapeutic actions. The same drawbacks hold for *DIAMON-1* [Steimann and Adlassnig, 1994a; Steimann and Adlassnig, 1994b], a two-stage monitoring system based on fuzzy sets.

RÉSUMÉ [Shahar and Musen, 1993; Shahar and Musen, 1996] performs temporal abstraction of time-stamped data without predefined trends. The system is based on a model of three basic temporal abstraction mechanisms: point temporal abstraction (a mechanism for abstracting the values of several parameters into a value of another parameter), temporal inference (a mechanism for inferring sound logical conclusions over a single interval or two meeting intervals) and temporal interpolation (a mechanism for bridging non-meeting temporal intervals). However, their approach is not applicable in our domain for the following reasons. First, *RÉSUMÉ* covers only limited domain dynamics (e.g., different classifiers for different degrees of parameters' abnormalities are not included). Second, it requires predefined domain knowledge to perform the temporal interpolation (e.g., gap functions) which is not available in our domain. Third, it concentrates on methods to cope with low-frequency observations which cannot easily be adapted for high-frequency domains concerning different properties of the domain. Fourth, the high level abstraction mechanism (pattern matching based on external and internal knowledge) is superfluous for therapy planning and not retrievable in our domain. Finally, *RÉSUMÉ* does not offer therapy planning.

Our approach benefits from using all available information about the state of the respiratory system of the patient based on temporal ontologies (time points and intervals [Allen, 1991; Dean and McDermott, 1987]), on different granularities (continuously and discontinuously assessed data) and on various kinds of data (quantitative and qualitative data). Our temporal data abstraction mechanism covers the different degrees of parameters' abnormalities and their corresponding dynamics (e.g., "the higher the degree of a parameter's abnormality the higher the positive change in the parameter's value which is classified as normal") as well as expected qualitative trend descriptions (e.g., "the transcutaneous partial pressure of oxygen ($P_{tc}O_2$) value should reach the normal region within approximately 10 to 20 minutes") to arrive at unified qualitative descriptions of parameters.

Our underlying concept of temporal data abstraction is expressed in schemata for data-point transformation and curve fitting. The schemata for data-point transformation represent regions of interest and attainable goals which are based on a single parameter's observations. The dynamically derived schemata for curve fitting represent attainable goals, dynamically derived thresholds to classify parameters, and clinically relevant expectations which are possible and commonly expressed in clinical routine during observation periods.

To keep our qualitative descriptions stable we apply smoothing and adjustment mechanisms. Additionally, we do not predefine absolute, time-dependent expected normal behavior of parameters during the whole observation process (like, Haimowitz and Kohane [Haimowitz, 1994; Haimowitz et al., 1995]), because the course of a parameter according to an absolute temporal dimension (axis) is not known in advance. We derive schemata for curve fitting in relation to the specific states of each

parameter. Improving or worsening of these parameters are best described as exponential functions. The costs for comparing such exponential functions are reduced by stepwise linearization. Finally, our approach is able to handle incomplete knowledge by applying repair methods or simplified system models.

5.3 VIE-VENT'S SYSTEM ARCHITECTURE

VIE-VENT is an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants [Miksch et al., 1993]. We have incorporated alarming, monitoring, and therapy planning tasks within one system to overcome important limitations of existing systems (e.g., GUARDIAN [Hayes-Roth et al., 1992]; SIMON [Uckun et al., 1993; Uckun et al., 1992]; NeoGanesh/Ganesh [Dojat and Sayettat, 1994; Dojat and Sayettat, 1995]; VentPlan [Rutledge et al., 1993]). The data-driven architecture of VIE-VENT consists of five modules: data selection, data validation, data abstraction, data interpretation and therapy planning. All these steps are involved in each cycle of data collection from monitors. VIE-VENT is especially designed for practical use under real-time constraints at neonatal ICUs (NICUs) and its various components are built in analogy to the clinical reasoning process.

VIE-VENT addresses the following central aims:

- (1) system architecture:
 - composing a modular, extensible system architecture with reusable components according to the main tasks of the monitoring and therapy planning process
 - implementing a real-time, open-loop system
- (2) data interpretation:
 - using all clinically relevant and available data (e.g., non invasive continuous blood gas measurements, qualitative clinical descriptions such as the degree of chest wall expansion)
 - reducing the number of parameters to context-sensitive and context-relevant parameters depending on the state of the neonate's respiratory system
 - dealing with faulty, vague, and missing data
- (3) knowledge-based monitoring and therapy planning:
 - implementing practical clinical and textbook knowledge according to the clinical reasoning process
 - performing monitoring and therapy planning in the absence of an appropriate structure-function model
 - using all available data: point, interval, trend data, and expected qualitative trend descriptions to arrive at unified qualitative descriptions of parameters which are used in the system model for data validation, data interpretation, and therapy planning
 - monitoring the patient during the whole artificial ventilation process (an initial phase, a phase of controlled ventilation (intermittent positive pressure ventilation, IPPV), a phase of weaning (intermittent mandatory ventilation, IMV) and a phase of returning to spontaneous breathing)
 - supporting individual strategies of mechanical ventilation

Figure 5.1 gives an overview of VIE-VENT's architecture. The left-hand side of Figure 5.1 illustrates the input data, namely transcutaneous blood gases, ventilator settings, invasive blood gas analyses, and clinical parameters. The labels of the arrows give the frequency of the measurements (e.g., "every 10 seconds" means, that the variables are available every 10 seconds and these variables are continuously assessed ones). The available knowledge is described in section 5.4. The right-hand side of Figure 5.1 expresses VIE-VENT's output. The system output consists mainly of therapeutic recommendations for changing the ventilator settings. Additionally, VIE-VENT gives warnings in critical situations, as well as comments and explanations about the state of the respiratory system of the newborn infant. The frequency of the output data depends on the state of the neonate's respiratory system and temporal preconditions according to the therapy planning strategies. The gray box in the middle of Figure 5.1 illustrates the interactions of the five modules of VIE-VENT. Neonatal respiratory function in our system model is represented by two processes, namely ventilation (CO₂ elimination) and oxygenation (oxygen uptake). The data selection module filters out context-relevant data for further analysis. The data validation and abstraction modules consist of two intertwined elements. First, the data validation uses the numerical values of the parameters to arrive at reliable values and the data abstraction transforms these validated values into unified qualitative descriptions based on different temporal ontologies. Second, the derived qualitative values are used again for data validation which may result in a revision of previously derived qualitative values. The data interpretation module classifies the state of the respiratory system of the newborn infant based on the unified qualitative values of the parameters. The therapy planning module formulates therapeutic actions based on the interpretation of monitoring data, prunes therapeutic actions, and verifies whether the therapeutic actions are effective.

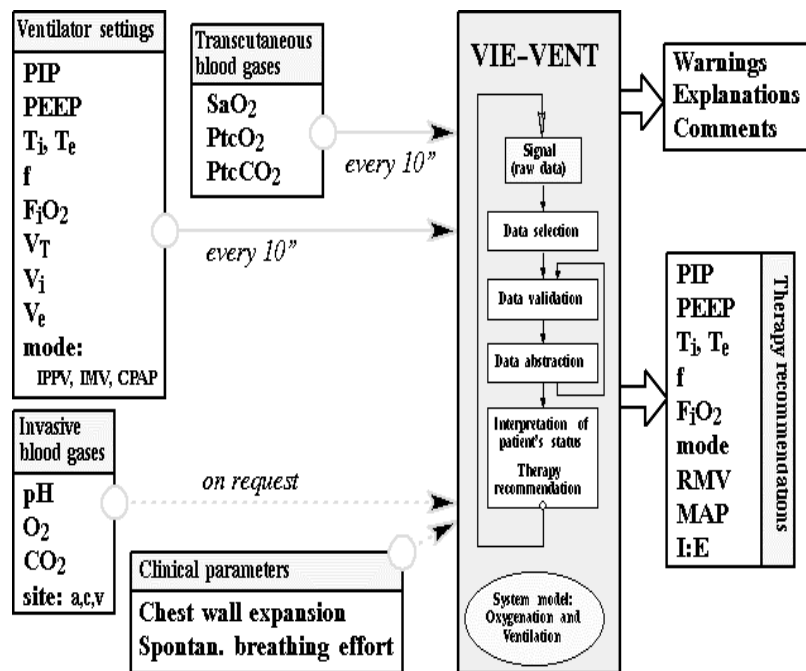


Figure 5.1 System architecture of VIE-VENT. The left-hand side shows the input data. The labels of the arrows give the frequency of the measurements. The right-hand side expresses VIE-VENT's output.

VIE-VENT is implemented in Clips [Clips, 1993; Giarratano and Riley, 1994], an expert system building tool integrating forward chaining rules and object-oriented programming. Most examples below are given in Clips notation.

5.4 THE AVAILABLE KNOWLEDGE

The available knowledge about the problem domain covers three main areas: observed data, classification knowledge, and temporal dynamic knowledge. In this section we describe the general idea of each area.

5.4.1 Observed Data

The observed data are the input data referring to the patient. According to the technical equipment of modern ICUs a huge volume of on-line data is available. Additionally, discontinuously assessed data and qualitative observations are available and needed for a global picture of the patient's condition and for an effective reasoning process. The data are divided according to their observed frequency and regularity as well as to their data types. We distinguish three kinds of data: continuously assessed quantitative data, discontinuously assessed quantitative, and qualitative data.

Continuously assessed input data are taken from the output of the data selection module every 10 seconds. Discontinuously assessed data are entered on request to the system depending on different conditions (e.g., critical ventilatory condition of the neonate, elapsed time intervals). VIE-VENT uses the following input parameters:

- (1) *continuous* quantitative data:
ventilator settings: F_iO_2 , f , PIP, PEEP, t_i , t_e , v_i , v_e , V_T
mode of ventilation: IPPV, IMV, CPAP
transcutaneous blood gases: $P_{tc}O_2$, $P_{tc}CO_2$, S_aO_2
- (2) *discontinuous* quantitative data:
clinical parameters (e.g., weight, age)
invasive blood gases: pH, PO_2 , PCO_2
site of blood gas measurements: arterial, capillary, venous
- (3) *discontinuous* qualitative data:
neonate's personal description (e.g., name, sex)
clinical parameters (e.g., chest wall expansion, spontaneous breathing effort)

5.4.2 Classification Knowledge

The classification knowledge represents time-point oriented concepts (e.g., the transformation of numerical data into qualitative regions of interest, functional and causal dependencies) and time-independent concepts (e.g., reliability ranking (invasive blood gases are more reliable, thus more important, than transcutaneous blood gases)).

5.4.3 Temporal Dynamic Knowledge

The temporal dynamic knowledge describes the relation among propositions of parameters over time (e.g., persistence of valid or plausible values of parameters during a time period, expected qualitative trend descriptions of parameters ("the parameter $P_{tc}O_2$ is moving one qualitative step towards the target range within 10 to 30 minutes")).

5.5 TEMPORAL DATA ABSTRACTION

The aim of the data abstraction process is to arrive at unified qualitative descriptions. The data abstraction is based on time points, time intervals and expected qualitative trend descriptions. Therefore, the data abstraction process has to handle various kinds of data occurring at different frequencies. Derived qualitative descriptions are used in the system model for data validation, data interpretation, and therapy planning. An advantage of using qualitative values is their uniform usability in the system model, irrespective of their origin. Adaptation to specific situations can easily be done by using specific transformation tables without changing the model of data interpretation and therapy planning. Additionally, by using qualitative values, an easily comprehensible and transparent system model can be developed.

VIE-VENT uses five different kinds of data abstraction:

- transformation of quantitative point data into qualitative values (schemata for data-point transformation)
- smoothing of data oscillating near thresholds
- smoothing of schemata for data-point transformation
- context-sensitive adjustment of qualitative values
- transformation of interval data (schemata for trend-curve fitting)

The schemata for data-point transformation convert time-stamped parameters' observations into qualitative values. To keep the qualitative values stable we apply various smoothing mechanisms. In critical states of the patient we have to adjust the qualitative values avoiding severe lung damages. The schemata for trend-curve fitting include the context-sensitive and expectation-guided temporal

knowledge to classify the observed parameters during time periods. The following subsections explain these methods.

5.5.1 Schema for Data-Point Transformation

The transformation of quantitative point data into qualitative values is usually performed by dividing the numerical value range of a parameter into regions of interest. Each region represents a qualitative value. It is comparable to the "point temporal abstraction"- task of Shahar and Musen [Shahar and Musen, 1993; Shahar and Musen, 1996].

The basis of the transformation of the blood gas measurements are *schemata for data-point transformation* relating single values to seven qualitative categories of blood gas abnormalities (qualitative *data-point* categories). The seven numerical regions of interest are not equal sized. The value range of an interval is smaller the nearer the target range. This is an important feature representing the dynamics related to the different degrees of parameters' abnormalities. It is extensively used in the schemata for trend-curve fitting (compare section 5.5.5). The schemata for data-point transformation are defined for all kinds of blood gas measurements depending on the blood gas sampling site (arterial, capillary, venous, transcutaneous) and the mode of ventilation (intermittent positive pressure ventilation (IPPV), intermittent mandatory ventilation (IMV), continuous positive airway pressure (CPAP)). The different modes of ventilation require specific predefined target values depending on different attainable goals. Figure 5.2 shows the schema of transcutaneous partial pressure of carbon dioxide ($P_{tc}CO_2$) during IPPV. For example, the transformation of the transcutaneous $P_{tc}CO_2$ value of 56 mmHg during IPPV results in a qualitative $P_{tc}CO_2$ value of *s2* ("substantially above target range") whereas during IMV it would represent a target value. IMV differs from IPPV in principle by the acceptability of higher CO_2 values (and analogous for pH, partial pressure of oxygen (O_2), and arterial oxygen saturation (S_aO_2) values). The $w_{i,x}$ values divide the qualitative regions. The transformation of interval data is based on these qualitative data-point categories which are described later.

The transformation of the data point values is stored in `qual_data_point_category` templates. The above mentioned example would be represented as

```
(qual_data_point_category
  (parameter PtcCO2)
  (mode_of_vent ippv)
  (site tc)
  (value s2))
```




Figure 5.2 Schema for data-point transformation of $P_{tc}CO_2$ during IPPV. On the left-hand side the abbreviations of the seven derived qualitative data-point categories are used. Their corresponding meaning is listed on the right-hand side. The gray region labeled by "normal" indicate the target range which should be achieved as soon as possible. The vertical gray box gives an example of the schema for data-point transformation of $P_{tc}CO_2$ during mode of ventilation IPPV. The labels $w_{i,x}$ indicates the thresholds of the regions of interest which are used in section 5.5.5.2. The [and] square brackets show the interval order (e.g.,]70, 148] is a left-side open, ride-side closed interval).

5.5.2 Data Smoothing

To avoid rapid changes of the qualitative categories triggered by data which oscillate near the thresholds of the schema for data-point transformation we apply a smoothing mechanism. The key idea is to keep the qualitative categories stable if the quantitative values cross the border to the next qualitative category just marginally for a few moments. Our smoothing mechanism is based on the size of the regions of interest, predefined ϵ regions and remaining time intervals. Alternative smoothing approaches could use statistical measurements (e.g., interval of confidence) or fuzzy sets to classify the parameter values.

The smoothing mechanism:

Let¹ a_t be the actual value at time t with $a_t \in [w_i, w_{i+1}]$, a_{t-1} be the value one time-step before (with $a_{t-1} \in [w_{i-1}, w_i]$ or $a_{t-1} \in [w_i, w_{i+1}]$ or $a_{t-1} \in [w_{i+1}, w_{i+2}]$), w_k be the borders of the qualitative data-point categories, $qual(a_m)$ be the related qualitative data-point categories at time point m , then

```
if    qual( $a_t$ )  $\neq$  qual( $a_{t-1}$ ) and (( $a_t \leq w_i + \epsilon$ ) or ( $a_t \geq w_{i+1} - \epsilon$ ))
then start smoothing
     $\forall a_m, m \in [t, t+x]$  :
        if qual( $a_m$ )  $\neq$  qual( $a_{t-1}$ ) and
            (( $a_m \leq w_i + \epsilon$ ) or
             ( $a_m \geq w_{i+1} - \epsilon$ ))           then qual( $a_m$ )  $\leftarrow$  qual( $a_{t-1}$ )
        else if (( $a_m > w_i + \epsilon$ ) or
                ( $a_m < w_{i+1} - \epsilon$ ))       then stop smoothing
with if | $w_{i+1} - w_i$ | > 3      then  $\epsilon = 2$ 
        else  $\epsilon = 1$ 
    and  $[t, t+x]$  be the remaining time interval
```

The smoothing mechanism starts if the current qualitative data-point category ($qual(a_t)$) is not equal to the previous qualitative data-point category ($qual(a_{t-1})$) and a_t is in the ϵ region. At the starting point, the actual qualitative category gets the value of the previous category. During the remaining time interval the new actual category $qual(a_m)$ gets the value of the category at the time point $t-1$ ($qual(a_{t-1})$) if the preconditions hold. The smoothing lasts as long as one of the following preconditions holds:

- (1) predefined time period (e.g., 5 minutes) since the start of smoothing (t) has not been elapsed and
- (2) a_m is in the ϵ region

Figure 5.3 gives an example of our smoothing mechanism. At time point t the smoothing mechanism is activated, because the two preconditions " $qual(a_t) \neq qual(a_{t-1})$ " and " $a_t \leq w_i + \epsilon$ " are satisfied. Therefore, the shifting of the qualitative categories starts at time point t . The gray arrows (\downarrow) illustrate the shifting of data values from the qualitative data-point category "s2" to the qualitative category "s1". At time point $t+3$ no shifting is necessary because the qualitative category is the same as at the starting point of the smoothing. The data smoothing lasts until time point $t+5$, because of failing of the second precondition at time point $t+6$. In this example, the predefined time period of 5 minutes has not been exhausted.

¹) In contrast to Figure 5.2, the second index (upper and lower region) has been eliminated to increase readability.




Figure 5.3 Example of data smoothing. The x axis represents the time, the data points are plotted every 30 seconds (the increments represent the next 30 seconds). The y axis represents the parameter value. The qualitative data-point categories are listed on the right-hand side and their corresponding thresholds on the left-hand side. The smoothing is activated at time point t and holds until time point $t+5$. The gray arrows (\downarrow) illustrate the shifting of data values from the qualitative data-point category "s2" to the qualitative category "s1".

5.5.3 Smoothing of Data-Point Transformation Schemata

The schemata for data-point transformation are defined for all modes of ventilation (IMV, IPPV) representing different target values. Changing the mode of ventilation would therefore result in an abrupt change of the schema for data-point transformation and by this in a sudden shift of the qualitative category. As a consequence, this could lead to recommendations for rather drastic changes of the ventilator settings. To avoid too abrupt changes of the qualitative categories, we smooth the thresholds of the schemata for data-point transformation within a predefined time period (three to eight hours depending on the "aggressiveness" of the user).

For example, if the mode of ventilation is changed from IPPV to IMV the thresholds of the schemata for data-point transformation are changed stepwise during eight hours in the case of a conservative user. This results in a slow change of the target range in the next eight hours, and with respect to the therapeutic consequences, in a graceful start of weaning process.

5.5.4 Context-Sensitive Adjustment of Qualitative Values

For extremely critical or life-threatening situations of a patient, the thresholds defined in the schemata for data-point transformation are too strict. In such cases we adjust the qualitative value of a parameter which is equal to a shift of the numerical threshold value. The adjustment of qualitative values holds as long as the precondition of "life-threatening situation" is true.

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 H₂O) is very high, higher P_{tc}CO₂ values are tolerated as better ones in order to prevent extreme pressure settings. The following two rules represent this kind of knowledge.

```
if (30 < PIP ≤ 35) and
    (PtcCO2 is "extremely below target range")
then (PtcCO2 is changed to "substantially below target range")
if (PIP > 35) and
    (PtcCO2 is "extremely below target range")
then (PtcCO2 is changed to "slightly below target range").
```

5.5.5 Transformation of Interval Data (Schema for Trend-Curve Fitting)

Similarly to the transformation of numerical data points to qualitative values, interval data are transformed to qualitative descriptions resulting in a verbal categorization of the change of parameters over time. Physicians' experiences about the expectations of a blood gas value changes over time in order to reach the target range in a physiologically proper way are expressed in verbal terms. For example, "the parameter P_{tc}O₂ is moving one qualitative step towards the target range within 10 to 30 minutes". These qualitative statements are called *expected qualitative trend descriptions*. The qualitative classification of the abnormality of a blood gas value resulted in different sized qualitative ranges (s3, s2, s1, normal, g1, g2, g3) as shown in section 5.5.1. Combining these qualitative data-point categories with the expected qualitative trend descriptions we obtain the *schemata for trend-curve fitting*. The schemata for trend-curve fitting express the dynamics of and the reactions to different degrees of parameters' abnormalities. A physician classifies a higher degree of a parameter's abnormality as more severe and a faster positive change of this parameter as normal. This is expressed in the different sizes of the data-point categories. The corresponding, dynamically derived, trends (depending on the expected qualitative trend descriptions) represent the different dynamic changes.

Depending on the particular purpose we vary the time period used in calculating the trends. We distinguish four kinds of trends based on our samples which are derived from new measurements every 10 seconds. The distinctions of the trends are guided by physiological criteria:

- (1) *very short-term* trend: sample of data points based on the *last* minute
- (2) *short-term* trend: sample of data points based on the *last 10* minutes
- (3) *medium-term* trend: sample of data points based on the *last 30* minutes
- (4) *long-term* trend: sample of data points based on the *last 3* hours

Table 5.1 summarizes the different purposes according to the different modules of monitoring and therapy planning (compare section 5.3). In addition, the first line of Table 5.1 shows the individual usage of the abstracted time-point data.

Table 5.1 Purposes of the different kinds of point data and trends. The rows indicate the different kinds of data and trends. The columns show the different purposes according to the various tasks of monitoring and therapy planning.

Kind of data	Purposes				
	Data Validation	Data Interpretation (State of the Respiratory System)	Therapy Planning		Assessment of Therapeutic Actions
			Changing Strategies	Therapeutic Actions	
Time-Point Data	√	√		√	
Very Short-Term Trend	√ (artifacts)	√			
Short-Term Trend	√	√		√	√
Medium-Term Trend		√		√	
Long-Term Trend		√	√		

The transformation of interval data into qualitative values is the last step of the data abstraction process. All necessary smoothing procedures are already done and only validated and therefore reliable data are involved. In case of missing or invalid measurements certain criteria of validity to proceed with the trend-based data abstraction process are needed.

In a monitoring process, the position of a measurement in the sequence of time-ordered data influences the reasoning process: namely, recent measurements are more important than historical measurements. Hence, criteria dealing only with an average distribution of measurements are insufficient. Due to this precondition we defined two criteria of validity to make sure, that the trend used is actually meaningful: a certain minimum amount of valid measurements within the whole time interval, and a certain amount of valid measurements within the last 20 percent of the time interval. These limits are defined by experts based on their clinical experience. They may easily be adapted to a specific clinical situation based on the frequency at which data values arrive. Table 5.2 shows the different conditions.

Table 5.2. Criteria of validity. The rows list the different kinds of trends. The columns list the validity criteria.

Kind of Trend	Valid Measurements <i>within the Whole Time Interval</i>	Valid Measurements <i>within the last 20 percent of the Time Interval</i>
Very Short-Term	at least 50 percent	100 percent
Short-Term	at least 40 percent	80 percent
Medium-Term	at least 30 percent	60 percent
Long-Term	at least 20 percent	40 percent

5.5.5.1 The Guiding Principle

The guiding principle of our approach is illustrated in Figure 5.4. The *schema for trend-curve fitting* transforms the different quantitative trend values (e.g., short-term or medium-term trends) into ten qualitative categories guided by physiological criteria. The x axis describes the discrete granularity of the representation in minutes. The y axis shows the $P_{tc}CO_2$ levels and the corresponding qualitative data-point categories. The value space of a parameter is divided into an upper and a lower region by the normal range. The dark gray area represents the expected qualitative trend description of a normal change of a parameter in the upper and the lower region, respectively. The derived qualitative trend categories are written in bold, capital letters.

Improving or worsening of parameters are modeled by exponential functions. An appropriate approach for classifying trend data is to transform the curve (borders of the dark gray area) shown in Figure 5.4 into an exponential function and to compare it with the actual growth rate. To classify the trend data we use a dynamic comparison algorithm which performs a stepwise linearization of the expected exponential function to overcome complexity (compare section 5.5.5.2).

For example, if a $P_{tc}CO_2$ data point is classified as *s1*, *s2* or *s3* ("... above target range") we would expect a therapeutic intervention to result in an decrease of type A2 (dark gray area) as "normal" trend.

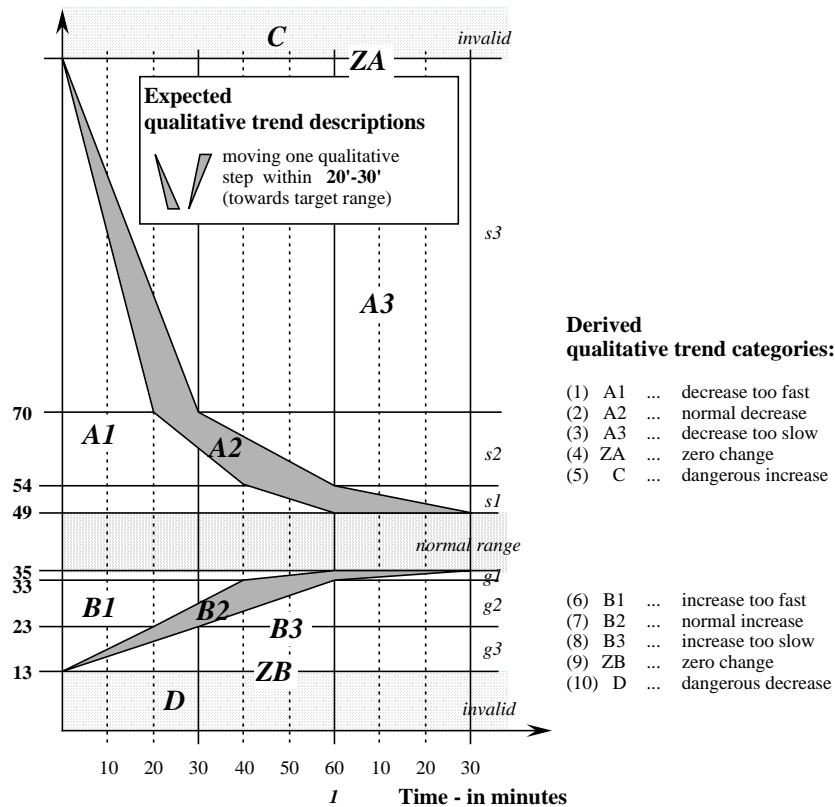


Figure 5.4 Schema for trend-curve fitting of $P_{tc}CO_2$. The x axis represents the time in minutes. The y axis represents the parameter value of $P_{tc}CO_2$ and the seven qualitative data-point categories. Their corresponding abbreviations and associated meaning are listed on the right-hand side. The light gray area indicates the normal (target) range. The dark gray area indicates the expected qualitative trend description of a normal change of a parameter in the upper and the lower region, respectively. The abbreviations of the ten derived qualitative trend categories are listed on the graphic.

5.5.5.2 The Dynamic Comparison Algorithm

The dynamic comparison algorithm classifies data within a time interval to a qualitative trend category depending on the relative position of corresponding data points and the expected qualitative trend descriptions. As an example, Figure 5.5 gives the schema for trend-curve fitting of $P_{tc}CO_2$ where we have reached a value of 85 mmHg after 58 minutes. The x axis describes the discrete granularity of the representation in minutes. The y axis shows the $P_{tc}CO_2$ levels. It indicates the quantitative values of data points (horizontal dotted lines indicate the thresholds). Their corresponding qualitative categories are listed on the right-hand side. Based on the guiding principle depicted in Figure 5.4, we compute the actual curve for selecting between the different qualitative categories. The striped area, A2, shows the expected normal development. The qualitative trend categories are written in bold, capital letters. They determine if an additional therapeutic action should be recommended (visualized with light-gray arrows in Figure 5.5)

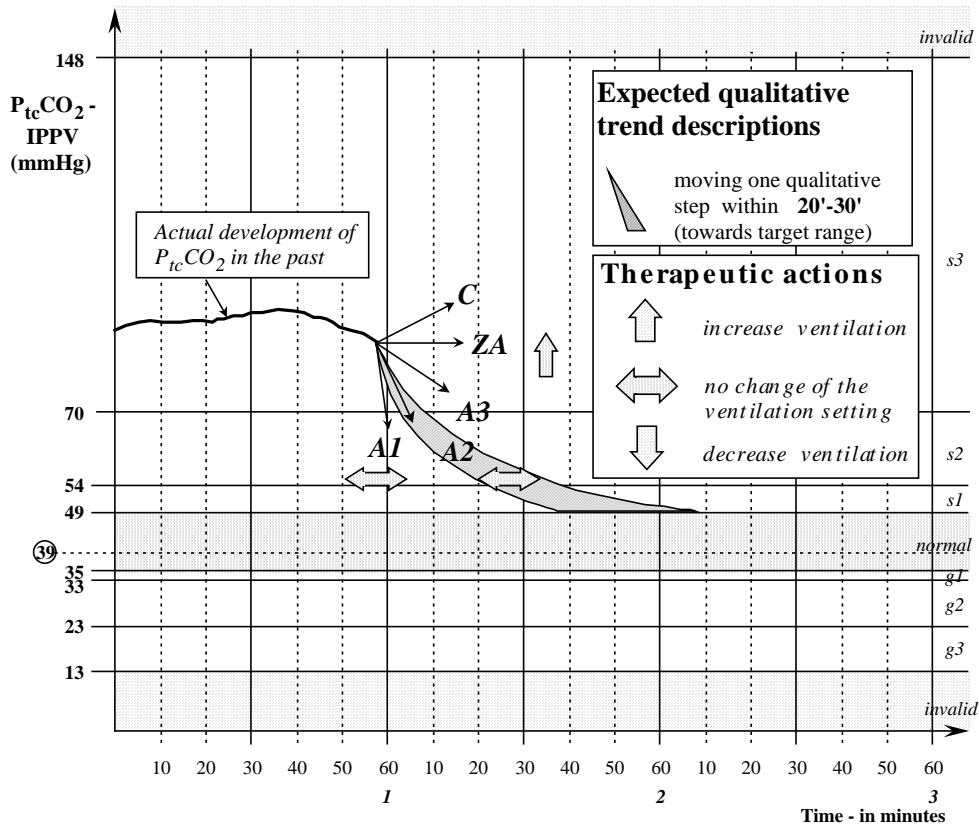


Figure 5.5 Example schema for trend-curve fitting of $P_{tc}CO_2$. The x axis represents time in minutes. The y axis shows the $P_{tc}CO_2$ levels. The $P_{tc}CO_2$ value has reached a value of 85 mmHg after 58 minutes. The striped area, A2, shows the expected normal development. The qualitative trend categories are written in bold, capital letters.

The growth rates are calculated and classified for all kinds of trends (very-short-, short-, medium-, and long-term). To increase readability, we only show the applicability of the method and not the results for all four kinds of trends. The algorithm is used for all trends.

The dynamic comparison algorithm consists of two steps:

- Step one:* Calculates the actual growth rate k_a using the linear regression model and two thresholds of the growth rate k_1 and k_2 depending on the relative position for the data points; k_1 and k_2 are used for discerning the qualitative trend categories A1, A2, and A3.
- Step two:* Classifies the qualitative trend category depending on the actual growth k_a , on the two thresholds k_1 , k_2 , and the qualitative region where the previous data point (a_{t-1}) belongs. In addition to k_1 and k_2 an ϵ range around zero is used to classify a trend as "ZA" and "ZB", respectively. The ϵ range is created on physiological grounds in order to support a wider range for defining "no change of a parameter".

The results of this algorithm are classifications of all parameters to one of the ten qualitative trend categories A1, A2, A3, ZA, B1, B2, B3, ZB, C, or D during a particular time period (explained in Figure 5.4). The value space of a parameter consists of an upper and a lower qualitative region divided by the target range. In the following we explain the algorithm used in the upper region. In this case we classify qualitative trend categories A1, A2, A3, ZA, or C. For the lower region the conditional elements (\leq , $>$) have to be changed to (\geq , $<$).

Step one:

First, we need to approximate the growth of the continuously assessed measurements $P_{tc}O_2$, $P_{tc}CO_2$ and S_aO_2 using a simple linear regression model ($E(Y) = a + k * X_p$, where $E(Y)$ is the expected value, X_i are the observed data points, a is a constant value (offset), and k is the growth rate). We assume, that the observations are mutually independent.

Choosing this simple linear regression model was influenced by practical clinical reasons: the only important characteristics of parameters used by physicians are on the one hand increases, decreases, or zero changes of parameters, and on the other hand too slow, too fast, or reasonable changes of parameters. Therefore, it would be superfluous to calculate a curve-fitting model of higher order with additional features for our application domain.

The calculation of the two thresholds of the growth rate k_1 and k_2 depends on the position of the data point. In the following we explain the calculation for the upper region.

Let²⁾ w_1 , w_2 , w_3 , and w_4 be the thresholds of the qualitative data-point categories above target range (as shown in Figure 5.2), where $\infty > w_1 > w_2 > w_3 > w_4 > 0$. Let a_t be the actual value at current time t , and $[t_{min}, t_{max}]$ the expected time interval for expected normal moving of one qualitative step towards the target range. The two thresholds of the growth rate k_1 and k_2 are then calculated as follows

²⁾ In contrast to Figure 5.2, the second index (upper and lower region) has been eliminated to increase readability.

$$\begin{aligned}
&\mathbf{if} \quad (a_t \leq w_3) \mathbf{then} \quad k_1 = \frac{(w_4 - w_3)}{t_{min}} \quad \mathbf{and} \quad k_2 = \frac{(w_4 - w_3)}{t_{max}} \\
&\mathbf{else} \quad \mathbf{if} \quad (a_t > w_2) \mathbf{then} \quad k_1 = \frac{\left(\frac{(w_2 - w_3)}{(w_1 - w_2)} (a_t - w_2) + w_3 - a_t \right)}{t_{min}} \quad \mathbf{and} \\
&\quad \quad \quad k_2 = \frac{\left(\frac{(w_2 - w_3)}{(w_1 - w_2)} (a_t - w_2) + w_3 - a_t \right)}{t_{max}} \\
&\mathbf{else} \quad k_1 = \frac{\left(\frac{(w_3 - w_4)}{(w_2 - w_3)} (a_t - w_3) + w_4 - a_t \right)}{t_{min}} \quad \mathbf{and} \\
&\quad \quad \quad k_2 = \frac{\left(\frac{(w_3 - w_4)}{(w_2 - w_3)} (a_t - w_3) + w_4 - a_t \right)}{t_{max}}
\end{aligned}$$

Step two:

Second, we classify the actual growth rate, k_a , as one of the ten qualitative trend categories. The previous qualitative data-point value is used to decide whether to apply the rules dealing with the region above or below the target value. The actual state is compared with the past development.

Let a_{t-1} be the value at previous time $t-1$, k_a be the actual growth rate, and k_1 and k_2 be the two thresholds calculated in step 1. The classification is as follows

```

if the data point  $a_{t-1}$  belongs to the region above the target value then
  if ( $k_a > \varepsilon$ ) then (qual_trend_category is C)
  else if ( $|k_a| \leq \varepsilon$ ) then (qual_trend_category is ZA)
    else if ( $k_a < k_1$ ) then (qual_trend_category is A1)
      else if ( $k_a < k_2$ ) then (qual_trend_category is A2)
        else (qual_trend_category is A3)

```

This process results in instantiations of qualitative trend descriptions for each blood gas measurement, for each kind of trend, and for each activated mode of ventilation. For subsequent explanations the qualitative values as well as the corresponding numerical values are stored in a qual_trend_category template.³⁾ E.g.,

```

(qual_trend_category
 (parameter PtcCO2) (kind_of_trend short) (mode_of_vent IPPV)
 (qual_trend_category A3)
 (numerical_growth -0.03) (numerical_const 87.6)
 (k1 -0.023) (k2 -0.016))

```

The temporal data abstraction process results in unified qualitative data-point and trend categories. These categories are extensively used in the data validation and the therapy planning processes. Therefore, the system model of these two processes is easily comprehensible and transparent. In the following section we show how the qualitative categories are applied to improve validation and therapy planning.

³⁾ A corresponding template specifies the expected qualitative trend description. The qualitative data-point categories of a_{t-1} and a_t are stored in similar templates (compare section 5.5.1).

5.6 DATA VALIDATION

The data abstraction mechanism described in the previous section is interwoven with the data validation process. First, the data validation process uses the numerical values of the parameters to arrive at reliable values which are transformed into unified qualitative descriptions by the data abstraction process. Second, it applies these derived qualitative descriptions to detect faulty measurements. The major aim of the data validation process is to detect faulty measurements or artifacts and finally to arrive at reliable measurements which may be used for further analysis tasks. An artifact is a situation where a measured variable does not reflect the clinical context. Undoubtedly, data validation is an important, but often neglected part of the monitoring and therapy planning process.

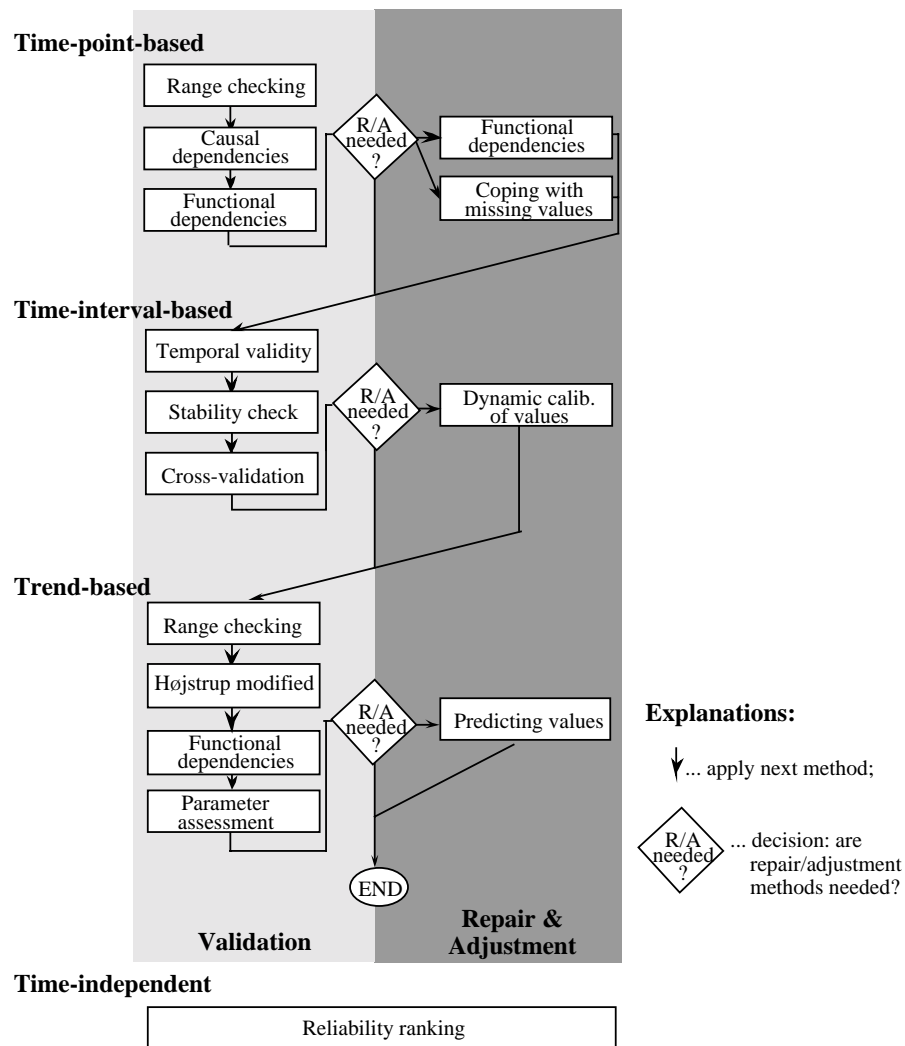


Figure 5.6 Overview and interaction of the components of the data validation and repair/adjustment modules. The left-hand side labels given indicate the different temporal-based ontologies. The left-hand side shows the validation methods. On the right-hand side the possible repair or adjustment methods are mentioned.

We perform a two-step data validation process based on different temporal ontologies: first, a context-sensitive examination of the plausibility of input data and second, applying repair and adjustment methods for correcting wrong or ambiguous data. The final result is a classification of the input data as "correct", "wrong", "unknown", or "adjusted". We divide our methods into four types based on their underlying temporal ontologies: time-point-based, time-interval-based, trend-based, and time-independent validation and repair. Figure 5.6 gives an overview of the particular categories and their interactions.

The data validation process uses numerical values as well as derived qualitative data-point and trend categories. The reasoning methods based on time points and time intervals represent a preprocessing for reasoning methods based on trends. Reasoning methods based on time points primarily perform static data validation which delivers the necessary preconditions to proceed with the trend-based validation. In this paper we concentrate on trend-based methods which utilize the qualitative categories. A detailed description of the whole process is given in [Miksch et al., 1996].

5.6.1 Assessment of Development of Parameters

The assessment procedure for the development of a particular parameter is based on the qualitative trend categories and the ordering of these categories. The preconditions to proceed are a positive judgment of the criteria of validity for calculating trends and a valid change allowed for a parameter (compare Figure 5.6). The simple method of comparing the actual numerical growth rate with a predefined growth threshold is not applicable in our case. The reason lies in the physiological behavior which results in different expected normal growth rates depending on the absolute values of the parameter. Therefore, we use the ten qualitative trend categories based on the short-term trends which reflect this dynamic aspect. The qualitative trend categories are divided by the target range into an upper and a lower region. According to these regions the ordering of the qualitative categories is defined as follows

ordering of the qualitative trend categories:

qualitative *upper* region: A1 - A2 - A3 - ZA - C

qualitative *lower* region: B1 - B2 - B3 - ZB - D

The assessment procedure compares the previous qualitative short-term trend category with the actual qualitative short-term trend category. If the actual category belongs to the same qualitative category or to a qualitative neighboring category of the previous category then the parameter is validated as "correct". Otherwise the parameter is classified as "wrong". We can only apply the short-term trends, because the very short-term trend reacts too rapidly to small oscillations of the values (compare Figure 5.7) and the medium- and long-term trends are too insensitive.

Figure 5.7 illustrates an example of the assessment procedure of the $P_{tc}CO_2$ time series. The derived qualitative trend and data-point categories are listed in the lower region of Figure 5.7. At 20:57 the $P_{tc}CO_2$ value is classified as "wrong", because the actual qualitative short-term trend category "A1" does not belong to "A3" (the same qualitative category as the previous value) or to "A2" or "ZA" (the neighboring categories of the previous category). The very short-term trend categories are not applicable for this purpose, because they show too rapid changes of the qualitative categories (e.g., at 20:45, 20:48, 20:49, 20:53, 20:55, and 20:57). The medium- and long-term trends react too slow, therefore they are not applicable in detecting faulty measurements.


```

(causal-explanation-validation
 (parameter PtcCO2)
 (classification adjusted)
 (repair-mechanism dynamic-calibration)
 (starting (day 941201) (time 13:23:00))
 (ending (day undefined) (time undefined)))

```

The reliability score of the adjustment/repair mechanism is predefined. For example, "dynamic-calibration" is a highly accurate mechanism, but "predicting values" is less accurate.

5.7 THERAPY PLANNING

After completion of the data validation and data abstraction process, the next two steps are interpretation of the respiratory status of the patient and therapy planning. Here we concentrate on the therapy planning process. The status interpretation task is described in detail in [Miksch et al., 1993]. The therapy planning module performs three tasks: formulating of therapeutic actions based on the status interpretation of monitoring data, pruning of therapeutic actions, and verifying the effectiveness of therapeutic actions.

5.7.1 Formulation of Therapy Recommendations

The first step in therapy planning is the collection of therapeutic recommendations which are possible in the current situation of the newborn infant. The qualitative temporal abstraction of monitoring variables makes it easy to use simple rules to activate therapeutic actions. The therapeutic actions and the corresponding qualitative trend categories are visualized with light-gray arrows in Figure 5.5. Rule R8-therapeutic-actions gives an example of such a rule.

```

Example: Rule R8-therapeutic-actions
(defrule R8-therapeutic-actions
  "activate-therapeutic-action-PtcCO2-ventilation"
  (phase (kind therapy_recommendation))
  (ventilation_phase (kind ippv))
  ?f1 <- (thp_recommendation ventilation)
  (qual_trend_category (parameter PtcCO2)
                       (kind_of_trend short)
                       (qual_trend A3|ZA|C))
  (qual_data_point_category (parameter PtcCO2)
                             (site tc) (value s1, s2, s3))
  (causal-explanation-validation (parameter PtcCO2)
                               (classification correct))
=>
  (retract ?f1)
  (assert (action (reason ventilation) (BG PtcCO2) (kind inc-f))
          (action (reason ventilation) (BG PtcCO2) (kind inc-pip))))

```

The essential preconditions for triggering therapeutic actions depend on the qualitative trend category of the short-term trend (expressed as `qual_trend_category` in the rule R8-therapeutic-actions) and the qualitative data-point category (expressed as `qual_data_point_category` in the rule R8-therapeutic-actions). If the qualitative data-point category is not normal and the qualitative trend category is A3 or ZA or C, then therapeutic actions are recommended (increase ventilator settings). This recommendation corresponds to the increase in ventilation shown by light-gray arrows in Figure 5.5. Quantity and frequency of an action depend on the degree of abnormality of the blood gas measurement (e.g., `s3` is worse than `s2`, therefore a larger amount of change is recommended) and the strategy of ventilation (e.g., aggressive or conservative). These features have been described in more detail in [Miksch et al., 1993]. The second fact `ventilation_phase` in the left-hand side (LHS) of rule R8-therapeutic-actions refers to

the mode of ventilation (i.e., IPPV) and indicates, that this rule belongs to the set of rules dealing with this phase of ventilation⁴). The last fact, *causal-explanation-validation*, delivers the necessary explanations of the data validation process, namely the classification of the particular validated parameter (compare section 5.6). The right-hand side (RHS) of rule *R8-therapeutic-actions* specifies the therapeutic actions. Each action-fact includes the kind of the recommended action and an explanation of the circumstances: the fact (*reason ventilation*) refers to "ventilation" process depending on the system model of ventilation⁵), (*BG PtcCO2*) refers to the relevant parameter, namely the blood gas measurement, and (*kind ?x*) determines which particular action has to take place (e.g., (*kind inc-pip*) means, that an increase of the peak inspiratory pressure (PIP) is recommended).

5.7.2 "Therapy Evaluator": Pruning of Therapy Recommendations

The process "formulation of therapy recommendations" creates a list of possible therapeutic actions according to the independently observed monitoring parameters. Such a list is given in Example 1.

Example 1:

```
((action (reason ventilation) (BG PtcCO2) (kind inc-f))
 (action (reason ventilation) (BG PtcCO2) (kind inc-pip))
 (action (reason oxygenation) (BG PtcO2) (kind dec-fiO2))
 (action (reason oxygenation) (BG PtcO2) (kind dec-pip))
 (action (reason oxygenation) (BG PtcO2) (kind dec-ti))).
```

As can be seen in Example 1, it is possible, that according to the blood gas measurements, contradictory (e.g., *dec-pip* and *inc-pip*) and/or too many therapeutic actions are recommended. Therefore, the duties of the "therapy evaluator" are to rank and to prune therapeutic actions. VIE-VENT distinguishes three different kinds of therapy evaluation: importance ranking of parameters, priority lists of attainable goals, and pruning of contradictory therapy recommendations. These methods are applied in the above-mentioned order.

Our approach differs from the existing mechanisms for assessing and comparing costs and benefits of therapeutic actions (e.g., the decision-theoretic approach in VentPlan [Farr and Fagan, 1989]) by using heuristic knowledge to rank and to prune important or unimportant therapeutic actions.

5.7.2.1 Importance Ranking of Parameters

Importance ranking⁶) of parameters specifies which therapeutic actions should take place first, according to the qualitative categories of the parameters. VIE-VENT uses the following rules:

- (1) Invasive blood gases are more reliable, thus more important, than transcutaneous blood gases.
- (2) S_aO_2 is more reliable than $P_{tc}O_2$.
- (3) Therapeutic actions depending on the qualitative data-point category " ... below target range" (*s1*, *s2*, *s3*) are more important than therapeutic actions depending on the qualitative data-point category " ... above target range" (*g1*, *g2*, *g3*).

⁴) The set of rules is divided into rules which hold for all phases of ventilation and rules which hold only for a specific phase. The phase of ventilation is related to the mode of ventilation as mentioned in section 5.3.

⁵) VIE-VENT's system model of ventilation is divided into two processes: oxygenation and ventilation. Different parameters are involved in these two processes (compare [Miksch et al., 1996]).

⁶) In the data validation component a similar method is used to validate the input data as reliable, called *reliability ranking*.

- (4) Within the range of the qualitative data-point category " ... below target range" (s1, s2, s3), therapeutic actions aiming to improve PO₂ or S_aO₂ are more important than therapeutic actions depending on PCO₂.

5.7.2.2 Priority lists of attainable goals

The priority lists of attainable goals specify the order in which the parameters should be changed. In particular, the priority lists indicate which parameter should reach which value space next. The value range of the parameters is divided into valid and invalid regions. Additionally, the valid region is divided into subintervals and each subinterval indicates a specific attainable value space. The first parameter in the priority list is chosen according to its value space. If a signal is received, that the previous therapeutic actions had twice resulted in an insufficient improvement of the patient's condition (compare section 5.7.3), the alternative parameter or the next parameter is chosen. The goals are separately defined for the blood gas measurements and ventilator settings.

The global goal for the blood gas measurements is to arrive at the normal range as soon as possible according to physiological criteria. The normal ranges of the different blood gas measurements are determined in the schemata for data-point transformation. They are different depending on the sampling site and the mode of ventilation, thus resulting in a different target range for each parameter (e.g., (PO₂, (transcutaneous, IMV), 40, 55), (PCO₂, (arterial, IPPV), 49, 35)).

The global concept of priority lists of the ventilator settings consists of a *general order of parameters* and a *particular priority list of attainable value spaces* which has to be reached first depending on the direction of change, the number of available parameters, and the relation of available parameters expressed in corresponding intervals of value space.

In VIE-VENT the attainable goals for the ventilator settings depend on the two processes of the system model: ventilation and oxygenation. The two ventilator settings peak inspiratory pressure (PIP) and frequency (f) as well as the optional information about the chest wall expansion are involved in the ventilation process. The *general order* for increasing is try "f before PIP" and the *general order* for decreasing is try "PIP before f". If the value of chest wall expansion, a discontinuously assessed parameter, is available then this additional precondition is added (e.g., if chest wall expansion is normal then increase f before PIP, if chest wall expansion is small then increase PIP before f). The *particular priority list* of attainable value spaces for PIP and f is listed in Table 5.3. The particular priority list gives values of equal priority of PIP and f in the same row. This list is used first. In case PIP and f have equal priority, the general order rule is used.

For example, we would like to decrease the ventilation. Let PIP be 28, f be 45, and chest wall expansion be normal, then PIP will stepwise decrease to 15 before starting to decrease f to 40.

Table 5.3 The particular priority list of attainable value spaces for PIP and f (ventilation process). The shadow field indicates the lower and upper plausible values of the parameters.

PIP (cm H ₂ O)	f (breaths/minute)
10	20
15	40
20	60
25	80
30	100
35	120
40	135
45	150

In the case of oxygenation a similar *general order* for increasing and decreasing and a *particular priority list* of attainable value spaces for the ventilator settings inspiratory oxygen concentration (F_iO_2), PIP, positive endexpiratory pressure (PEEP), and inflation time (t_i) are defined to rank the therapy recommendations.

5.7.2.3 Pruning of Contradictory Therapy Recommendations

The pruning of contradictory therapy recommendations identifies the important and consistent recommendations from a set of conflicting therapy recommendations. Currently, we apply a straight forward strategy to deal with contradictory therapy recommendations because most of the possible occurrences are already handled by the other features of the "therapy evaluator" which are applied first.

HEURISTIC:

If an increase as well as a decrease of the same parameter is recommended then delete both therapeutic actions.

In Example 1 (see above), both actions (*inc-pip*) and (*dec-pip*) are deleted.

5.7.3 Assessment of Therapeutic Actions

The duty of the assessment of therapeutic actions is to assess their therapeutic efficiency. VIE-VENT defines a therapeutic action as "effective" or "ineffective" based on the ordering of the qualitative trend categories and a delay time.

(1) *ordering of the qualitative trend categories:*

qualitative *upper* region: C ZA A3 A2 and A1 A2

qualitative *lower* region: D ZB B3 B2 and B1 B2

In the qualitative upper region the best attainable category is A2. A2 is the qualitative trend category which represents the physiologically best (normal) trend of a parameter. There are two possible ways to reach the category A2: first, it is possible to get there from the most severe (worst) category C to ZA, then to A3 and finally A2. Second, we may get there from category A1.

The same concept is used in the qualitative lower region for categories D, ZB, B3, B2, and B1 respectively.

(2) *delay time:*

The delay time is based on experience, namely we need to wait a certain amount of time until we are able to see clearly the outcome of therapeutic actions in terms of monitored transcutaneous blood gas values. The delay time for all observed blood gas parameters is 10 minutes.

If a therapeutic action has taken place, we would expect the qualitative trend category to improve by one step in the ordering direction after the delay time. If no qualitative improvement can be detected then the action is assessed as "ineffective" and the book-keeping procedure is activated. The book-keeping procedure collects and counts all previously ineffective therapeutic actions. If the book-keeping procedure verifies, that a previous therapeutic action has *twice* failed to improve the patient's clinical condition, it forwards a signal to the "therapy evaluator" (namely the component dealing with the priority lists of attainable goals, discussed in section 5.7.2.2) to trigger an alternative action.

5.8 DATA VISUALIZATION: A GLOBAL, COMPREHENSIVE PICTURE

Physicians and nurses have to observe multiple channels of different kinds of data at the same time. Additionally, they have to classify the observed data over a time period. Inevitably, they may have difficulties tracking all these data to detect severe situations and/or significant trends. Therefore, it is essential to display the case-relevant, interpreted data and the interweaving of data as comprehensibly as possible [Gershon, 1994].

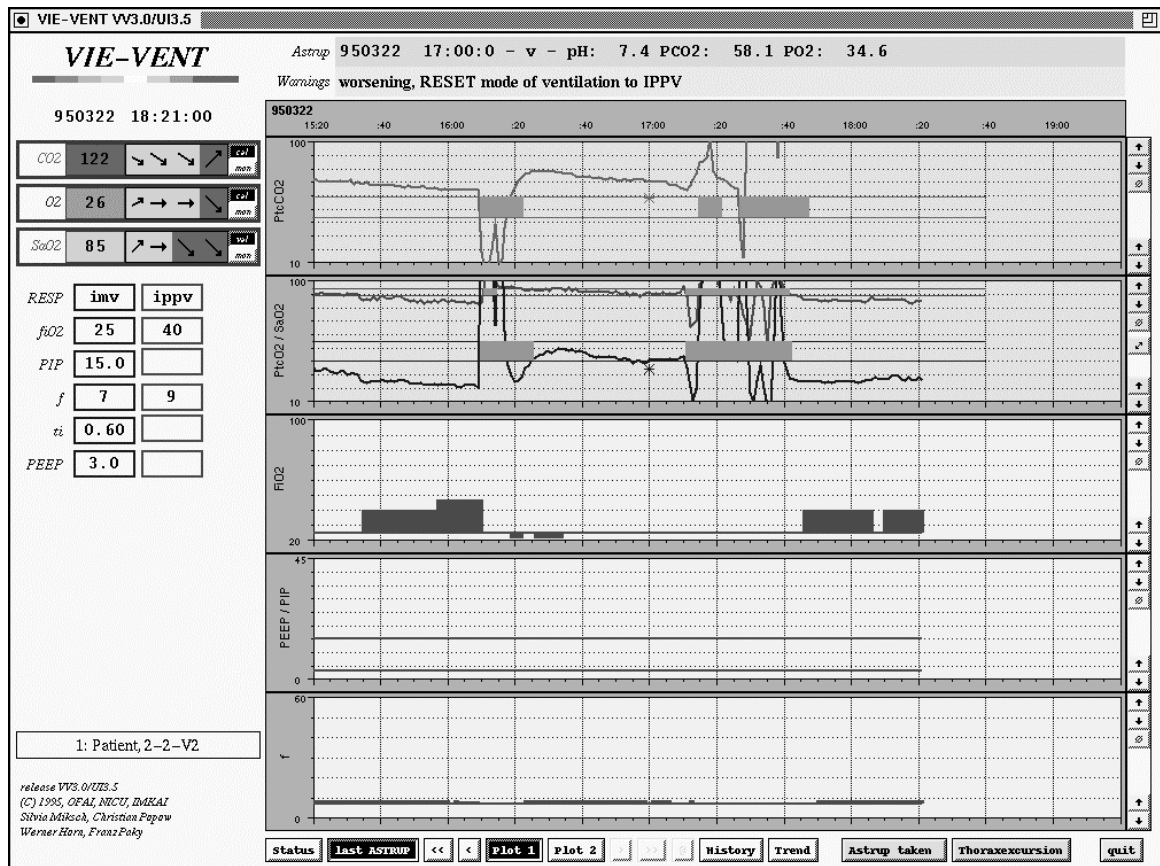


Figure 5.8 The user interface of VIE-VENT. The left-hand side region shows the blood gas measurements, their corresponding qualitative temporal abstraction on the top and the actual and recommended ventilator settings below. The right-hand side region gives plots of the most important parameters over the last four hours, namely transcutaneously assessed blood gas measurements and some ventilator settings.

Figure 5.8 shows the user interface of VIE-VENT. The left-hand side shows the blood gas measurements and their corresponding qualitative temporal abstraction on the top. The actual ventilator settings (first row, e.g., F_iO_2 is 25), and VIE-VENT's therapeutic recommendations at the current time (second row, e.g., increase F_iO_2 to 40) is given below. The upper right-hand side shows two status lines. First, the time and the values of the last invasive blood gas measurements (labeled by "Astrup") are given. The invasive assessed blood gas measurements are also displayed as colored asterisks in the corresponding transcutaneous blood gas curve. Second, additional warnings are labeled by "Warnings" (e.g., "worsening" means, that VIE-VENT detected, that the respiratory system of the neonate is worsening). The lower right-hand side gives plots of the most important parameters over the last four hours. Scrolling to previous time periods is possible by pressing the buttons (<<) for a four-hour step backward, (<) for an one-hour step backward, (>>) for a four-hour step forward, or (>) for an one-hour step forward, respectively. Additional information and explanations about other parameters, the history, and the temporal abstraction can be retrieved on users' request (pressing the buttons <Plot 2>, <History> or <Trend>, respectively). The temporal abstraction of the blood gas measurements is displayed in the left-hand upper corner: the qualitative data-point categories are expressed using color chart with different gradation (e.g., deep pink represents values extremely above the target range (s3), salmon represents values substantially above the target range (s2), lime green represents values extremely below the target range (g3)). The qualitative trend categories are visualized by colored arrows. The four arrows show the directions of the very-short, short, medium,

and long-term trends. The therapeutic recommendations are displayed as red vertical lines in the corresponding curve of the ventilator setting. Important future features are explained in the next section.

5.9 PRACTICAL USEFULNESS - FIRST CLINICAL EXPERIENCES

The design process of VIE-VENT concentrates on cyclic-interweaving tasks including phases of knowledge acquisition, system implementation, evaluation, and refinements. After several technical and subject evaluation steps we ended up with our current version of VIE-VENT. The realization of an on-line data acquisition enabled the observation of "real patients" instead of simulated patients as in earlier steps of our system development. Therefore, we have tested VIE-VENT in real clinical scenarios. We have particularly evaluated the usefulness and applicability of VIE-VENT's data validation, data abstraction, therapy planning, and data visualization components in real-clinical environments. Below we illustrate the practical usefulness based on our first experiences with on-line acquired patients.

Observing real patient records shows the importance of a secure and trustable data validation component. The two upper plots in Figure 5.8 shows the trajectories of the three most important, continuously observed measurements, namely $P_{tc}CO_2$ in the first upper plot as well as $P_{tc}O_2$ and S_aO_2 in the second upper plot. The gray areas between the two vertical lines (which represent the desired normal region of the blood gas measurement) express invalid measurements during a time period. The huge number of measurements classified as invalid is quite astonishing although it reflects the real-clinical environments. Invalidities are due to particular external events (e.g., calibration of transcutaneous sensors every three to four hours, endotracheal suctioning of secretions) that have to be performed regularly, on-line transmission problems, and unexplainable circumstances.

The considerable oscillation detected in the S_aO_2 values, sometimes more than 5 units (measurement in %), came as a great surprise. Usually the clinical staff observes the parameters on a digital display, where only spot data are displayed. Therefore, the clinical staff can hard observe the development of the S_aO_2 values during longer time periods and figure out the particular shapes of S_aO_2 values. These oscillations will be subject to future analyses. Important questions could be to combine the oscillation of the S_aO_2 values during different modes of ventilation as well as the comparison of conventional artificial ventilation with high-frequency artificial ventilation.

The usefulness of the qualitative trend categories and their visualizations has been manifested in different ways. First, they enable the physicians to get a closer insight into their medical reasoning process. They result in a refinement of our therapy recommendation component as well as in an adjustment and an enhancement of our priority lists and reliability ranking of parameters. Second, the qualitative trend categories improved the data validations component. Third, applying the qualitative trend categories for formulation and assessment of therapeutic actions results in a graceful weaning process avoiding too abrupt changes of therapeutic recommendations. In Figure 5.8 the therapeutic recommendations are displayed as red vertical lines in the corresponding curve of the ventilator setting. The therapeutic recommendations show a very consistent and reasonable picture, except in cases where the curves of the blood gas measurements express an unexplainable shape. Our data validation methods could not cope with these unexplainable shapes and failed. These findings result in a fine tuning of data validation method which should be completed during our current phase of evaluation.

5.10 DISCUSSION

In developing our methods for temporal data abstraction and applying our derived qualitative descriptions, we were motivated by monitoring and therapy planning in real-world environments and our application domain (artificially ventilated newborn infants at NICUs). But our approach is not limited to the presented domain. For applying our approach to other domains, the given problems

must have certain characteristics. These requirements are mainly derived by the real-world orientation and by the available properties of data and knowledge.

First, our approach is applicable in domains where a vast volume of high-frequency, on-line data and additional quantitative and qualitative off-line data are available. Data can be delivered at various frequencies (e.g., high or low frequency data), at various regularities (e.g., continuously or discontinuously assessed data), and be of various types (e.g., quantitative or qualitative data). All available data and knowledge should be included in the approach accordingly. Second, the observed data are more faulty than expected, but reliable data are needed for the reasoning processes. Third, the underlying structure-function models of the domain are poorly understood, because incomplete knowledge is involved. Therefore, well-known theories on data analysis are not applicable. But domain experts can express their experiences and knowledge with verbal descriptions which can be structured. Fourth, the main concentration lies in monitoring and therapy planning instead of diagnosis. Fifth, the final aim is to develop a real-time, open-loop system which has to act and to react in time-constraint situations immediately. It should be a data-driven approach which can cope with time-critical situations. Therefore, the required time to proceed with sophisticated reasoning methods is seldom available. Finally, a visualization of the observed data and the derived actions during different time periods are needed.

On one hand the most common methods of interpreting continuously assessed data (e.g., time-series analysis, probabilistic classifiers) are insufficient in dynamically changing environments where the temporal dimension covering the course of a parameter and the interdependencies of different parameters over time has to be taken into account. On the other hand we could not simply apply the knowledge-based approaches from e.g., *TrenDx* [Haimowitz, 1994; Haimowitz et al., 1995] or *RÉSUMÉ* [Shahar and Musen, 1993; Shahar and Musen, 1996], because of the lack of orientation of real environment, of predefined trend descriptions, of missing therapy planning, etc.

5.11 CONCLUSION

We demonstrated a very powerful temporal abstraction mechanism which combines all available information about the patient's state of the respiratory system to perform the temporal abstraction process. In designing our abstraction mechanism, we have concentrated on knowledge-based monitoring and therapy planning in real-clinical environments. First, the characteristics of the data and of the quality control have to be taken into account: vast volumes of data available to the medical staff, occurring at various frequencies, of various regularities and types, unexpected high volume of faulty data (artifacts), increased demands on quality control and assurance. Second, the problem definitions are not as clear as expected, because the underlying structure-function models for predicting the time course of clinical parameters are poorly understood and incomplete knowledge and complexity are involved. Therefore, theories on data analysis are not applicable.

Our methods are able to handle incomplete domain knowledge. The lack of an appropriate, reliable, structure-function model is substituted by a combination of data-point categories and expected qualitative trend descriptions. These schemata for trend-curve fitting include the dynamics and reactions to the different degrees of parameters' abnormalities. The results of our temporal abstraction process are unified qualitative descriptions which incorporate knowledge about point and interval data, as well as expected qualitative trend descriptions. These qualitative descriptions are suitable for improving the data validation, the interpretation of the patients' state of the respiratory system, the formulation of therapeutic actions, as well as for verifying whether these actions are effective within a certain period. Moreover, we apply a smoothing and adjustment mechanism to keep the qualitative descriptions stable and to react to critical states.

The first clinical experiences show, that the enhancement of our temporal data validation and abstraction methods has improved our therapy planning component remarkably. They guarantee a graceful weaning process avoiding too abrupt changes of parameters.

Acknowledgments

The current phase of the project is supported by the "Jubiläumsfonds der Oesterreichischen Nationalbank", Vienna, Austria, project number 4666 and by "Erwin Schrödinger Auslandstipendium, Fonds zur Förderung der wissenschaftlichen Forschung", J01042-MAT. We greatly appreciate the support given to the Austrian Research Institute of Artificial Intelligence (ÖFAI) by the Austrian Federal Ministry of Science, Research, and the Arts, Vienna.

References

- [Allen, 1991] Allen, J. F. (1991). Time and Time Again: The Many Ways to Represent Time. *International Journal of Intelligent Systems*, 6:341-55.
- [Avent and Charlton, 1990] Avent, R. K. and Charlton, J. D. (1990). A Critical Review of Trend-Detection Methodologies for Biomedical Monitoring Systems. *Critical Reviews in Biomedical Engineering*, 17(6):621-59.
- [Barahona and Christensen, 1994] Barahona, P. and Christensen, J. P., editors (1994). *Knowledge and Decisions in Health Telematics, The Next Decade*. Amsterdam: IOS Press.
- [Clips, 1993] Clips (1993). *Clips Reference Manual, Version 6.0*, Software Technology Branch. Lyndon B. Johnson Space Center, NASA.
- [Dean and McDermott, 1987] Dean, T. L. and McDermott, D. V. (1987). Temporal Data Base Management. *Artificial Intelligence*, 32(1):1-55.
- [Dojat and Sayettat, 1994] Dojat, M. and Sayettat, C. (1994). Aggregation and Forgetting: Two Key Mechanisms for Across-Time Reasoning in Patient Monitoring. In Kohane, I. S., Uckun, S., Coiera, E., Patil, R., and Stefanelli, M., editors, *Proceedings of the AI in Medicine: Interpreting Clinical Data*, pages 33-36, AAAI Press, Menlo Park.
- [Dojat and Sayettat, 1995] Dojat, M. and Sayettat, C. (1995). A Realistic Model for Temporal Reasoning in Real-Time Patient Monitoring. *Applied Artificial Intelligence*, 10(2):121-143.
- [EURISIC, 1995] EURISIC (1995). *European User's Requirements for Information Systems for Intensive Care*. European Society of Intensive Care Medicine. Working Group for Technical Developments. Kuopio, Finland.
- [Farr and Fagan, 1989] Farr, B. R. and Fagan, L. M. (1989). Decision-Theoretic Evaluation of Therapy Plans. In Kingsland, L. C., editor, *Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-89)*, pages 188-192, IEEE Computer Society Press, Washington D.C.
- [Gershon, 1994] Gershon, N. (1994). From Perception to Visualization. In Rosenblum, L., Earnshaw, R. A., Encarnacao, J., Hagen, H., Kaufman, A., Klimenko, S., Nielson, G., Post, F., and Thalmann, D., editors, *Scientific Visualization - Advances and Challenges*, pp. 129-139, Academic Press: London.
- [Giarratano and Riley, 1994] Giarratano, J. and Riley, G. (1994). *Expert Systems - Principles and Programming*. Second Ed. PWS Publishing Company, Boston.
- [Goebel and Schader, 1992] Goebel, H. and Schader, M. (1992). *Datenanalyse, Klassifikation und Informationsverarbeitung*. Physica, Heidelberg.
- [Goldsmith and Karotkin, 1993] Goldsmith, J. P. and Karotkin, E. H. (1993). *Assisted Ventilation of the Neonates*. Third Ed. Saunders, Philadelphia.
- [Haimowitz, 1994] Haimowitz, I. J. (1994). *Knowledge-Based Trend Detection and Diagnosis*. Ph.D. Thesis, Massachusetts Institute of Technology.

- [Haimowitz et al., 1995] Haimowitz, I. J., Le, P. P., and Kohane, I. S. (1995). Clinical Monitoring Using Regression-Based Trend Templates. *Artificial Intelligence in Medicine*, 7(6):473-496.
- [Hammond, 1990] Hammond, K. J. (1990). Explaining and Repairing Plans That Fail. *Artificial Intelligence*, 45(1-2):173-228.
- [Hayes-Roth et al., 1992] Hayes-Roth, B., Washington, R., Ash, D., Hewett, R., Collinot, A., Vina, A., and Seiver, A. (1992). GUARDIAN: A Prototype Intelligent Agent for Intensive-Care Monitoring. *Artificial Intelligence in Medicine*, 4(2):165-185.
- [Kay, 1993] Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing*. PTR Prentice Hall, Englewood Cliffs, New Jersey.
- [Miksch et al., 1996] Miksch, S., Horn, W., Egghart, G., Popow, C., and Paky, F. (1996). Monitoring and Therapy Planning without Effective Data Validation are Ineffective. In Kohane, I. S., Patil, R., Shahar, Y., Szolovits, P., and Uckun, S., editors, *Proceedings of the Artificial Intelligence in Medicine: Applications of Current Technologies*, AAAI Spring Symposium Series, pages 119-123, AAAI Press, Menlo Park, CA.
- [Miksch et al., 1993] Miksch, S., Horn, W., Popow, C., and Paky, F. (1993). VIE-VENT: Knowledge-Based Monitoring and Therapy Planning of the Artificial Ventilation of Newborn Infants. In Andreassen, S., Engelbrecht, E., and Wyatt, J., editors, *Proceedings of the Artificial Intelligence in Medicine, 4th Conference on Artificial Intelligence in Medicine Europe (AIME-93)*, pages 218-129, IOS Press, Amsterdam.
- [Perlman et al., 1995] Perlman, M., Claris, O., Hoa, Y., Pandit, P., Whyte, H., Chipman, M., and Lui, P. (1995). Secular Changes in the Outcomes to Eighteen to Twenty-Four Months of Age of Extremely Low Birth Weight Infants with Adjustment for Changes in Risk Factors and Severity of Illness. *Journal of Pediatrics*, 126:75-87.
- [Rutledge et al., 1993] Rutledge, G. W., Thomsen, G. E., Farr, B. R., Tovar, M. A., Polaschek, J. X., Beinlich, I. A., Sheiner, L. B., and Fagan, L. M. (1993). The Design and Implementation of a Ventilator-Management Advisor. *Artificial Intelligence in Medicine*, 5(1):67-82.
- [Shahar and Musen, 1993] Shahar, Y. and Musen, M. A. (1993). RÉSUMÉ: A Temporal-Abstraction System for Patient Monitoring. *Computers and Biomedical Research*, 26(3):255-273.
- [Shahar and Musen, 1996] Shahar, Y. and Musen, M. A. (1996). Knowledge-Based Temporal Abstraction in Clinical Domains. *Artificial Intelligence in Medicine, Special Issue Temporal Reasoning in Medicine*, 8(3):267-298.
- [Stefanelli, 1992] Stefanelli, M. (1992). Therapy Planning and Monitoring. *Artificial Intelligence in Medicine*, 4(2):189-190.
- [Steimann and Adlassnig, 1994a] Steimann, F. and Adlassnig, K.-P. (1994a). Clinical Monitoring with Fuzzy Automata. *Fuzzy Sets and Systems*, 61:37-42.
- [Steimann and Adlassnig, 1994b] Steimann, F. and Adlassnig, K.-P. (1994b). Two-Stage Interpretation of ICU Data Based Fuzzy Sets. In Kohane, I. S., and et al., editors, *Proceedings of the AI in Medicine, Interpreting Clinical Data*, pages 152-156, AAAI Press, Menlo Park, CA.
- [Uckun et al., 1993] Uckun, S., Dawant, B. M., and Lindstrom, D. P. (1993). Model-based Diagnosis in Intensive Care Monitoring: the YAQ Approach. *Artificial Intelligence in Medicine*, 5(1):31-48.
- [Uckun et al., 1992] Uckun, S., Dawant, B. M., Manders, E. J., and Lindstrom, D. P. (1992). SIMON: An Integrated Approach to Patient Monitoring in Critical Environments. In Lun, K. C., and al., e., editors, *Proceedings of the MEDINFO-92*, pages 564-569, North-Holland, Amsterdam.