in: *Proceedings of 5th Conference on Artificial Intelligence in Medicine Europe (AIME 95), June 25 - 28, 1995, Pavia, Italy*

Therapy Planning Using Qualitative Trend Descriptions

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
E-mail: silvia@ai.univie.ac.at

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

[3] NICU, Division of Neonatology, Department of Pediatrics, 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 a method of therapy planning applicable in the absence of an appropriate curve-fitting model. It incorporates knowledge about data points, data intervals, and expected qualitative trend description to arrive at unified qualitative descriptions of parameters (temporal data abstraction). Our approach benefits from derived qualitative values which can be used for recommending therapeutic actions as well as for assessing the effectiveness of these actions within a certain period. It results in an easily comprehensible and transparent concept of therapy planning. Furthermore, we improved the system model of data interpretation and therapy planning by using importance ranking of variables, priority lists of attainable goals, and pruning of contradictory therapy recommendations.*

Our methods are applicable in domains where an appropriate curve-fitting model is not available in advance. We have applied them in the field of artificial ventilation of newborn infants. The utility of our approach is illustrated by VIE-VENT, an open-loop knowledge-based system for artificially ventilated newborn infants.

Keywords:
Knowledge-based Monitoring and Therapy Planning (Temporal Reasoning), Artificial Ventilation, Newborn Infants

1 Introduction: The Needs For Special Therapy Planning Concepts

The care of critically ill patients in intensive care units (ICUs) is complex, involving interpretation of many variables, comparative evaluation of therapy options, and control of patient-management parameters. The technical improvement of the ICUs' equipment makes a huge amount of data available to the medical staff, and even skilled physicians frequently suffer from this information overload. Additionally, there are increased demands on the quality control and quality assurance ([8]: EURISIC-European User Requirements for Intensive Care).

Knowledge-based systems for intelligent alarming, diagnosis, monitoring and therapy planning have separately been developed. In contrast to diagnosis, which tries to find the best explanation for the actual situation of a patient, monitoring and therapy planning imply actions: *monitoring* indicates observing the course of a patient's condition under a given therapy, and assessing whether the selected therapeutic action is effective and the predicted improvement of the patient's condition occurs. *Therapy planning* involves selecting which therapeutic actions may improve the patient's condition, predicting the outcome, and adopting a therapeutic plan according to some explicitly defined preferences on the predicted condition of the patient [14]. Recent studies pointed out the challenge and the need to integrate all these activities within a unique framework, especially when dealing with dynamic domains ([1], [2]).

Control theory or statistical analysis is useful to allow for a straightforward mapping of monitoring data and appropriate therapeutic actions. In view of the lack of appropriate curve-fitting models for predicting the time course of clinical variables, as in the case of artificial ventilation, these methods are unsuitable. However, it is possible to express expected trends, like "the $P_{tc}O_2$ value should reach the normal region in approximately 10 to 20 minutes". We therefore tried to overcome the missing of an appropriate curve-fitting model by applying a dynamic temporal data abstraction mechanism based on spot and trend data analysis as well as expected qualitative trend descriptions to arrive at unified qualitative descriptions of variables. These qualitative values are used in the system model of data interpretation and therapy planning. An advantage of using qualitative values is their unified usability in the system model, no matter of which origin they are.

In the first part of this paper we introduce the system architecture of VIE-VENT. In the second part we focus on important components of the temporal data abstraction process to arrive at unified qualitative descriptions of data *points* and *trend*s. In the third part we explain VIE-VENT's therapy planning module based on the interpretation of the patient's health condition, on pruning of therapeutic actions, and on verifying whether therapeutic actions are effective.

2 VIE-VENT's System Architecture

Developing VIE-VENT, an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants [10], we incorporated alarming, monitoring, and therapy planning tasks in one system to overcome some of the limitations of existing systems. The data-driven architecture of VIE-VENT consists of several modules: data selection, data validation, data abstraction, data interpretation and therapy planning. All these steps are involved in a single cycle of data collection from monitors. VIE-VENT is especially designed for practical use under real-time constraints at neonatal ICUs (NICUs) and the various components are built in analogy to the clinical reasoning process.

On the one hand VIE-VENT's whole input data set can be divided into continuous and discontinuous data. Continuous data (e.g., blood gas measurements, like $P_{tc}O_2$, S_aO_2, $P_{tc}CO_2$, and ventilator settings, like PIP, F_iO_2) are taken from the output of the data selection module every 10 seconds. Discontinuous data are entered into the system on request by the user depending on different conditions (e.g., critical ventilatory condition of the neonate, elapsed time intervals, missing monitoring data). The system output consists in primarily therapeutic recommendations for changing the ventilator setting. Additionally, VIE-VENT gives warnings in critical situations, as well as comments and explanations about the health condition of the neonate.

On the other hand VIE-VENT's variables can be divided into dependent (e.g., blood gas measurements, chest wall expansion) and independent variables (e.g., ventilator settings). The therapeutic actions are composed of independent variables.

VIE-VENT monitors the patient during the whole artificial ventilation process. We divide the whole period into four phases: initial phase, controlled ventilation, weaning, returning to spontaneous breathing. Transition from one phase to the next is handled by rules depending on the amount of artificial ventilation needed.

3 Preconditions for Temporal Data Abstraction and Therapy Planning

The problem of planning artificial ventilation of newborn infants - as in other medical fields, like pediatric growth [6] - lies in the lack of an appropriate curve-fitting model to predict the development of physiological variables from actual measurements. Therefore our first effort was to approximate the growth of continuously assessed measurements such as $P_{tc}O_2$, $P_{tc}CO_2$, and S_aO_2 using a simple linear regression model ($E(Y) = a + k X_i$) 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 assumed that observations are mutually independent and have the same variance.

Choosing this simple linear regression model was influenced by practical clinical reasons: the only important characteristics of variables used by physicians for a real therapeutic action are on the one hand increases, decreases, or zero changes

of variables, and on the other hand too slow, too fast, or reasonable changes of variables. Therefore it would be superfluous to calculate a curve-fitting model of higher order with additional features for our purpose.

An appropriate curve-fitting function would be an exponential function with variables improving towards the normal range after a therapeutic action. On the basis of this concept we compare the actual trend with a stepwise linearized function representing the exponential curve. This decreases the complexity of a comparison of exponential functions and ensures responsiveness of the system. Additionally, the linear approximation is reasonable and applicable for small time intervals.

Based on physiological criteria, four kinds of trends of our 10 seconds data samples can be discerned:

(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.

Comparing different kinds of trends is a useful method of assessing the result of previous therapeutic actions, of detecting if oscillation is too rapid, and of isolating the occurrence of artifacts (compare [9]).

4 Data Abstraction

The aim of the data abstraction process is to arrive at unified qualitative descriptions of data points and trend data. It transforms quantitative measurements into qualitative values, which can be used in the system model for data interpretation and therapy planning. 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 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 data points into qualitative values, transformation of trend data, dynamic calibration of values, context-sensitive adjustment of qualitative values, and smoothing for data oscillating in the neighborhood of thresholds. The data abstraction process dealing with spot data (only data points are involved in contrast to data abstraction based on data intervals as well as combinations of data points and data intervals) has been described in [9]. We will focus on the first two components.

4.1 Transformation of Data Points (Data-Point-Transformation□Scheme)

The transformation of quantitative data points into qualitative values is usually performed by dividing the numerical range of a variable into regions of interest. Each region stands for a qualitative value. The region defines the only common property of the numerical and qualitative values. It is comparable to Shahar's et al. [13] "point temporal abstraction".

The basis of the transformation of the blood gas measurements are *data-point-transformation schemata* relating single values to seven qualitative categories of blood gas abnormalities (qualitative *data-point*-categories):

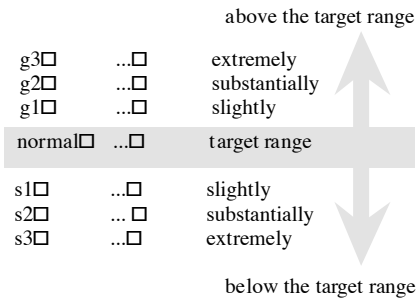

These data-point-transformation schemata are defined for all kinds of blood gas measurements depending on the blood gas sampling site (arterial, capillary, venous, transcutaneous) and the mode of ventilation (IPPV, IMV). The different modes of ventilation require specific predefined target values depending on different attainable goals. Fig. 1 shows the scheme of $P_{tc}O_2$ during IPPV. For example, the transformation of the transcutaneous $P_{tc}O_2$ value of 91 mmHg during IPPV results in a qualitative $P_{tc}O_2$ value of *g3* ("extremely above target range"). The $w_{i,x}$ -values divide the qualitative regions. The transformation of trends is based on these qualitative data-point-categories, which are described in the following section.

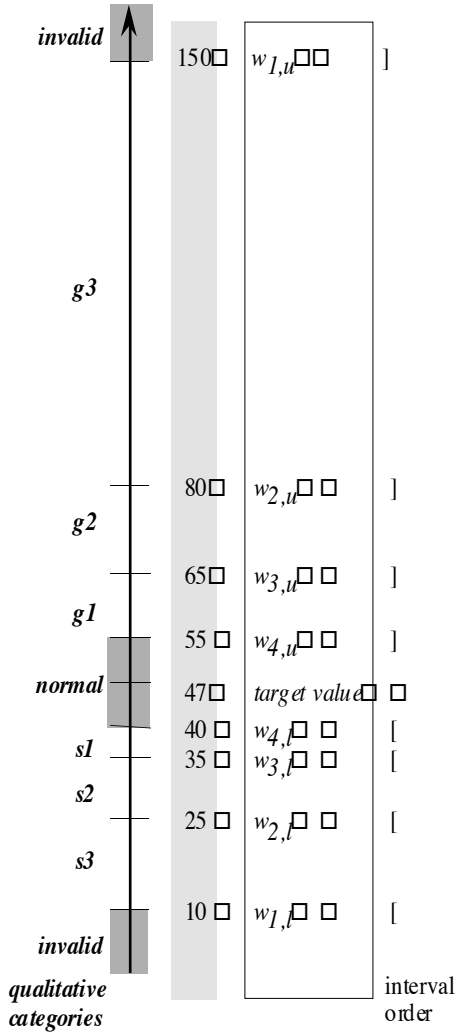

Fig. 1. Data-point-transformation scheme of $P_{tc}O_2$ during IPPV

4.2 Transformation of Trend Data (Trend-Curve-Fitting Scheme)

The transformation of trend data into qualitative values is based on the combination of qualitative data-point-categories and the qualitative descriptions of the expected behavior of a variable (*expected qualitative trend description*s; e.g., "variable $P_{tc}O_2$ is moving one qualitative step towards the target range within 10 to 30 minutes"). These *trend-curve-fitting schemata* transform the quantitative trend values into ten qualitative categories guided by physiological criteria (Fig. 2).

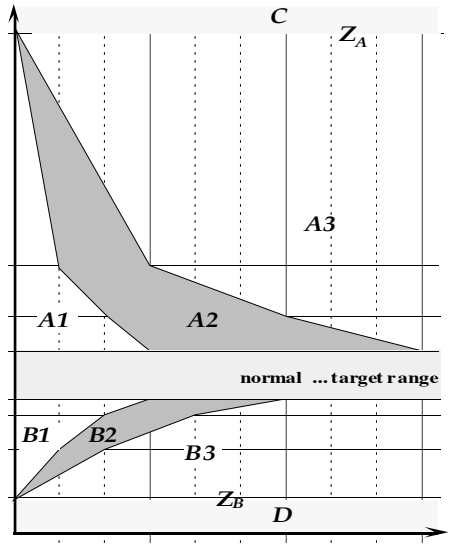

(1) *A1* ... decrease too fast
(2) *A2* ... normal decrease
(3) *A3* ... decrease too slow
(4) Z_A ... zero change
(5) *C* ... dangerous increase

(6) *B1* ... increase too fast
(7) *B2* ... normal increase
(8) *B3* ... increase too slow
(9) Z_B ... zero change
(10) *D* ... dangerous decrease

Fig. 2. Trend-curve-fitting scheme of $P_{tc}O_2$

A qualitative trend-category depends on the relative position of corresponding data points. For example, if a $P_{tc}O_2$ data point is classified as *g1, g2* or *g3* (" ... above target range") we would expect a therapeutic intervention to result in a decrease of type A2 as "normal" trend. If the data point lies in the target range ("*normal*") no therapeutic action is recommended.

As an example, Fig. 3 gives the trend-curve-fitting scheme of $P_{tc}O_2$ where we have reached a value of 91 mmHg after 43 minutes. The x-axis describes the discrete granularity of the representation in minutes. The y-axis shows the $P_{tc}O_2$ levels. It indicates the quantitative values of data points (at thresholds horizontal dotted lines are drawn) and their corresponding qualitative categories are listed on the right side. Based on the guiding principle depicted in Fig. 2, we compute the actual curve for selecting between the different qualitative categories. The stripped area A2 shows the expected normal development. The qualitative trend-categories are written in bold, capital letters and determine if an additional therapeutic action should be recommended (visualized with light-gray arrows in Fig. 3 and described in chapter 5).

An appropriate approach to classify trend data is to transform the curve (borders of the dark gray area) shown in Fig. 3 to an exponential function and to compare it with the actual growth rate. We used a dynamic comparison algorithm to classify the trend data, which performs a stepwise linearization of the expected exponential function to overcome complexity. It consists of two steps:

Step one: calculates the actual growth rate k_a using the linear regression model explained in chapter 3 and two thresholds of the growth rate k_1 and k_2 depending on the relative position of data point; k_1 and k_2 are used for discerning the qualitative trend-categories

Step two: classifies the qualitative trend-category depending on the actual growth k_a, the two thresholds k_1, k_2 and the qualitative region where the data point belongs. In addition to k_1 and k_2 we use an ε-range around zero to classify a trend as "Z_A" and "Z_B", respectively. The ε-range is created on physiological grounds in order to support a wider range for defining no change of a variable.

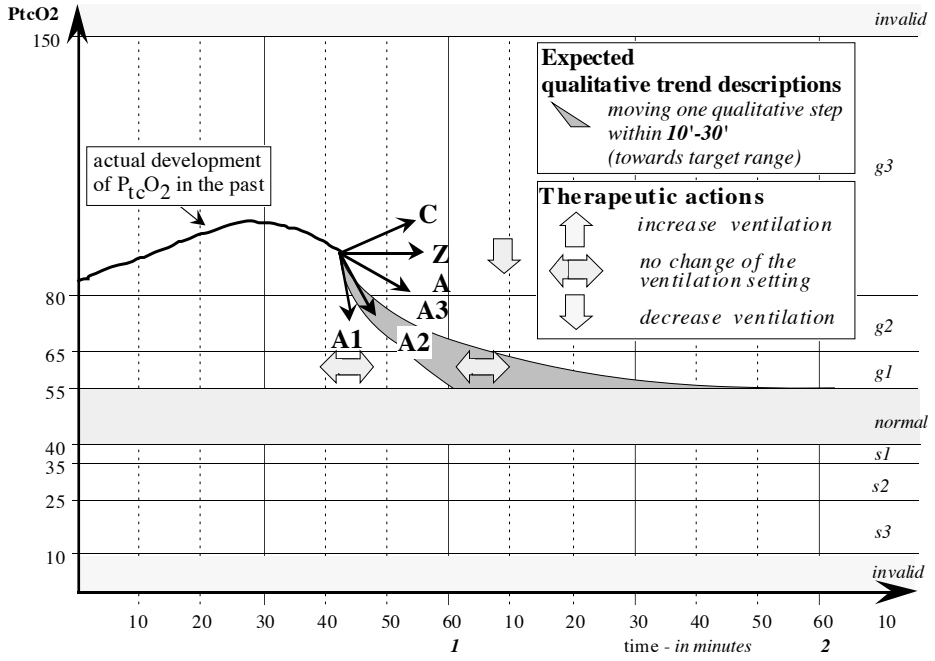

Fig. 3. Trend-curve-fitting scheme of $P_{tc}O_2$

Step one: The value space of a variable 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. For the lower region the conditional elements ($\leq$, $>$) have to be changed to ($\geq$, $<$).

Let[1] w_1, w_2, w_3, and w_4 be the thresholds of the qualitative data-point-categories above target range (as shown in Fig. 1), 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 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 in Fig. 4.

if $(a_t \leq w_3)$ **then**

$$k_1 = \frac{(w_4 - w_3)}{t_{min}} \text{ and } k_2 = \frac{(w_4 - w_3)}{t_{max}}$$

else if $(a_t > w_2)$ **then**

$$k_1 = \frac{\left(\frac{(w_2 - w_3)}{(w_1 - w_2)}(a_t - w_2) + w_3 - a_t\right)}{t_{min}} \text{ and}$$

$$k_2 = \frac{\left(\frac{(w_2 - w_3)}{(w_1 - w_2)}(a_t - w_2) + w_3 - a_t\right)}{t_{max}}$$

else

$$k_1 = \frac{\left(\frac{(w_3 - w_4)}{(w_2 - w_3)}(a_t - w_3) + w_4 - a_t\right)}{t_{min}} \text{ and}$$

$$k_2 = \frac{\left(\frac{(w_3 - w_4)}{(w_2 - w_3)}(a_t - w_3) + w_4 - a_t\right)}{t_{max}}$$

Fig. 4. Calculation of the two thresholds

[1] In contrast to figure 1, the second index has been eliminated to increase readability.

Step two:

if the data point 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)

The result of this process are instantiations of qualitative trend descriptions for each dependent variable (namely blood gas measurements) for each kind of trend, and for each mode of ventilation. For explanations the qualitative values as well as the corresponding numerical values are stored in a **`qualitative_trend`** template.[2] E.g.,

```
(qualitative_trend
     (BG_name PtcO2) (kind_of_trend very-short) (mode_of_vent IMV)
     (qualt_trend_category A3)
     (numerical_growth -0.03)(numerical_const 93.5)
     (k1 -0.39) (k2 -0.13))
```

5 Therapy Planning

The therapy planning module consists of formulation of therapeutic actions based on the interpretation of monitoring data, of pruning of therapeutic actions, and of verifying whether the therapeutic actions are effective. The data interpretation module has been described in [10].

5.1 Formulation of Therapy Recommendations

The qualitative abstraction of monitoring variables allows for creation of simple rules to activate therapeutic recommendations. Rule **`R5-therapeutic-actions`** (Fig. 5) gives an example of such a rule. The rule syntax is defined in Clips notation (v6.02, COSMIC/NASA), a forward chaining rule and/or object based development system.

The essential preconditions to trigger therapeutic actions depend on the qualitative trend-category (expressed as **`qualt_trend_category`** for a short-term trend) and the qualitative data-point-category (expressed as **`qual_data_point_category`**). If the qualitative data-point-category does not belong to the set {normal} and the qualitative trend-category belongs to the set {A3, ZA, C}, then an action is recommended (e.g., to decrease ventilator settings). Amount 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 [10]. The second fact **`ventilations_phase`** in the left-hand-side (LHS) of the rule refers to the mode of ventilation (i.e., IMV) and indi-

---

[2] A corresponding template specifies the expected qualitative trend description.

cates that this rule belongs to the set of rules dealing with this phase of ventilation.[3] The right-hand-side (RHS) of rule **R5-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 oxygenation)** refers to the depending process "oxygenation" of the system model of ventilation[4], **(BG ?BG)** refers to the depending variable, namely the blood gas measurement, and **(kind ?x)** determines which particular action has to take place (e.g., **(kind dec-pip)** means that a decrease of the peak inspiratory pressure (PIP) is recommended).

```
(defrule R5-therapeutic-actions
            "activate-therapeutic-action-PtcO2-oxygenation"
     (phase (kind therapy_recommendation))
     (ventilations_phase (mode imv))
     ?f1 <- (thp_recommendation oxygenation)
     (qualt_trend_category      (BG_name PtcO2)
                         (kind_of_trend short)
                         (qualt_trend_category A3|ZA|C))
     (qual_data_point_category (kind PtcO2) (value ~normal))
     =>
     (retract ?f1)
     (assert (action (reason oxygenation) (BG PtcO2) (kind dec-fiO2))
             (action (reason oxygenation) (BG PtcO2) (kind dec-peep))
             (action (reason oxygenation) (BG PtcO2) (kind dec-pip))
             (action (reason oxygenation) (BG PtcO2) (kind dec-ti))))
```

Fig. 5. Example: Rule R5-therapeutic-actions

5.2 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 $\prec$ ZA $\prec$ A3 $\prec$ A2 and A1 $\prec$ A2
qualitative *lower* region: D $\prec$ ZB $\prec$ B3 $\prec$ B2 and B1 $\prec$ B2

In the qualitative upper region the best attainable category is A2. 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, 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* for all observed dependent variables 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 proce-

---

[3] 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 chapter 2.

[4] VIE-VENT's system model of ventilation is divided into two processes: oxygenation and ventilation. Different parameters are involved in these two processes (compare [10]).

dure collects and counts all previously ineffective therapeutic actions. If the bookkeeping 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, compare chapter 5.3.2) to trigger an alternative action.

5.3 "Therapy Evaluator": Pruning of Therapy Recommendations

The process "formulation of therapy recommendations" renders a list of possible therapeutic actions according to the independently observed monitoring variables, e.g.,

```
Example 1:
      ((action (reason oxygenation) (BG PtcO2)(kind dec-fiO2))
       (action (reason oxygenation) (BG PtcO2)(kind dec-peep))
       (action (reason oxygenation) (BG PtcO2)(kind dec-pip))
       (action (reason oxygenation) (BG PtcO2)(kind dec-ti))
       (action (reason ventilation) (BG PtcCO2)(kind inc-pip))
       (action (reason ventilation) (BG PtcCO2)(kind inc-f))).
```

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 variables, priority lists of attainable goals, and pruning of contradictory therapy recommendations. These methods are applied in the order as they are mentioned in the following.

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 ([5], [12]) by using heuristic knowledge to rank and to prune important or unimportant therapeutic actions.

5.3.1 Importance ranking of Variables

Importance ranking[5] of variables specifies which therapeutic actions should take place first according to the qualitative categories of the dependent variable. VIE-VENT uses the following rules:

(1) invasive blood gases are more reliable, thus more important, than transcutaneous blood gases;

(2) SaO_2 is more important 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) within the range of the qualitative data-point category " ... below target range" (s1, s2, s3) therapeutic actions aiming to improve PO_2 or S_aO_2 are more important than therapeutic actions depending on PCO_2;

---

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

5.3.2 Priority Lists of Attainable Goals

The priority lists of attainable goals specify the order in which the variables should be changed. In particular the priority lists rank which variable should reach which value space next. The value range of the variables is divided into a valid and an invalid range. Additionally, the valid range is divided in subintervals and each subinterval indicates a specific attainable value space. The first variable 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 chapter 5.2), the alternative variable or the next variable is chosen. The goals are separately defined for dependent (e.g., blood gas measurement $P_{tc}CO_2$) and independent (e.g., ventilator setting PIP) variables.

The global goal for the dependent variables, namely 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 data-point-transformation schemata. It is different depending on the sampling site and the mode of ventilation. This results in a different target range for each variable (e.g., (PO_2, (transcutaneous, IMV), 40, 55), (PCO_2, (arterial, IPPV), 49, 35)).

The global concept of priority lists of the independent variables consists of a *general order of variables* and a *particular priority list of attainable value space*s which has to be reached first depending on the direction of change, the number of available variables, and the relation of available variables expressed in corresponding intervals of value space. In VIE-VENT the attainable goals for the independent variables, namely 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 *global order* for decreasing is try "PIP before f". If the value of chest wall expansion 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 1.

In the case of decreasing, PIP has to be first decreased to 45 (45 is the plausible upper limit of PIP), then f has to be decreased to 150 (150 is the plausible upper limit of f), then PIP has to be decreased stepwise to 40 then f has to be decreased stepwise to 120, and so on. For example, let PIP be 28, f be 45, and chest wall expansion be normal, then PIP will stepwise be decreased to 15 before starting to decrease f to 40. The actual amount of change depends on the degree of blood gas abnormality, namely the qualitative data-point-category. The priority list of attainable goals just determines which variable has to be taken first depending on the attainable intervals.

PIP (cm H_2O)	**f** (breaths/minute)
10	20
15	40
20	60
25	80
30	100
35	120
40	135
45	150

Table 1. The particular priority list of attainable value spaces for PIP and f.

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 F_iO_2, PIP, PEEP and T_i are defined to rank the therapy recommendations.

5.3.3 Pruning of Contradictory Therapy Recommendations

The pruning of contradictory therapy recommendations identifies the important 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 variable is recommended
then we delete both therapeutic actions.
```

In the example 1 (see above), both actions `(inc-pip)` and `(dec-pip)` are deleted.

6 Related Work

In the past decade several strategies have been developed to support monitoring and therapy planning. The drawback of most approaches is that they were developed for low-frequency data (like in pediatric growth monitoring where new data are collected several months or a year apart (e.g., TrenDx [6]), in diabetes where new data arrive three or four times a day (e.g., [11]), or in artificial ventilation management using only invasively determined variables (e.g., GUARDIAN [7]). Monitoring and therapy planning of high-frequency data of ICU patients require different strategies for data validation, monitoring and therapy planning.

A comparable, but closed-loop approach is NeoGanesh/Ganesh [3, 4]. It is a rule-based system designed for weaning of artificially ventilated adults. Its temporal model consists of aggregation and forgetting to choose an appropriate

treatment plan. Our temporal data abstraction covers and extends parts of RESUME's [13] temporal abstraction of time-stamped data. E.g., we extended the "point temporal abstraction" mechanism of RESUME with expected qualitative trend descriptions to arrive at unified qualitative values. The main purpose of RESUME is temporal data abstraction and trend detection. It lacks on activating therapeutic actions and assessing the benefits of therapeutic actions. The same drawback holds for DIAMON-1 [15], a two-stage monitoring system based on fuzzy sets.

7 Conclusion

We demonstrate a method to improve monitoring and therapy planning in the absence of an appropriate a priori curve-fitting model. Our approach benefits from derived qualitative values of data points, data intervals and expected qualitative trend descriptions (temporal data abstraction). These qualitative values can be used for recommending therapeutic actions as well as for verifying whether these actions are effective within a certain period. Additionally, we improved the system model of data interpretation and therapy planning by using importance ranking of variables, priority lists of attainable goals, and pruning of contradictory therapy recommendations. Our approach results in an easily comprehensible concept of therapy planning.

Acknowledgment

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

References

1. Barahona P., Christensen J.P.(eds.): Knowledge and Decisions in Health Telematics, The Next Decade, IOS, Amsterdam, 1994.
2. Console L., Molino G., Torasso P.: Some New Challenges for Artificial Intelligence in Medicine, in Barahona P., Christensen J.P.(eds.), *Knowledge and Decisions in Health Telematics*, IOS, Amsterdam, 1994.
3. Dojat M., Brochard L., Lemaire E., Harf A.: A Knowledge-Based System for Assisted Ventilation of Patients in Intensive Care Units, *International Journal of Clinical Monitoring and Computing,* 9, pp.239-50, 1992.
4. Dojat M., Sayettat C.: Aggregation and Forgetting: Two Key Mechanisms for Across-Time Reasoning in Patient Monitoring, in Kohane I.S., et al.(eds.), *AI in Medicine: Interpreting Clinical Data,* AAAI Press, Menlo Park, pp.33-36, 1994.
5. Farr B.R., Fagan L.M., Decision-theoretic Evaluation of Therapy Plans, in Kingsland L.C.(ed.), *Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-89),* IEEE Computer Society Press, Washington D.C., pp. 188-92, 1989.
6. 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.

7. 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.

8. Kari A.: Quality Control and Quality Assurance in Finland, in Metnitz P.G.H.(ed.), *Patientdaten Management System auf Intensivstationen*, Workshop Notes, Wiener Intensivmedizinische Tage (WIT 94), Vienna, 1994.

9. Miksch S., Horn W., Popow C., Paky F.: Context-Sensitive Data Validation and Data Abstraction for Knowledge-Based Monitoring, in Cohn A.G. (ed.), *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI 94)*, Wiley, Chichester, UK, p. 48-52, 1994.

10. 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 of the 4th Conference on Artificial Intelligence in Medicine Europe (AIME-93)*, IOS Press, Amsterdam, pp.218-29, 1993.

11. Ramoni, M., Riva A., Stefanelli M., Patel V.L.: Forecasting Glucose Concentration in Diabetic Patients Using Ignorant Belief Networks, in Kohane I.S., et al.(eds.), *AI in Medicine: Interpreting Clinical Data,* AAAI Press, Menlo Park, pp.33-36, 1994.

12. 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.

13. Shahar Y., Tu S.W., Musen M.A.: Knowledge Acquisition for Temporal Abstraction Mechanisms, Special Issue: Knowledge Acquisition for Therapy-Planning Tasks, *Knowledge Acquisition*, 4(2), 1992.

14. Stefanelli, M.: Therapy Planning and Monitoring, *Artificial Intelligence in Medicine*, 4 (2), pp. 189-90, 1992.

15. Steimann F., Adlassnig K.-P.: Two-Stage Interpretation of ICU Data Based Fuzzy Sets, in Kohane I.S., et al.(eds.), *AI in Medicine: Interpreting Clinical Data*, AAAI Press, Menlo Park, pp.152-6, 1994.