

Context-Sensitive and Expectation-Guided Temporal Abstraction of High-Frequency Data

Silvia Miksch^{1)*}, Werner Horn^{1,2)}, Christian Popow³⁾, Franz Paky⁴⁾

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

²⁾ Department of Medical Cybernetics and Artificial Intelligence, University of Vienna
Freyung 6, A-1010 Vienna, Austria
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
email: christian.popow@vm.akh-wien.ac.at

⁴⁾ Department of Pediatrics, Hospital of Mödling
Weyprechtgasse 12, A-2340 Mödling, Austria

*) Currently, visiting scholar at Knowledge System Laboratory (KSL), Stanford University, CA, USA.

Abstract

Therapy planning benefits from derived qualitative values or patterns which can be used for recommending therapeutic actions as well as for assessing the effectiveness of these actions within a certain period. Dealing with high-frequency data, shifting contexts, and different expectations of the development of parameters requires particular temporal abstraction methods to arrive at unified qualitative values or patterns.

This paper addresses context-sensitive and expectation-guided temporal abstraction methods. They incorporate knowledge about data points, data intervals, and expected qualitative trend patterns to arrive at unified qualitative descriptions of parameters (temporal data abstraction). Our methods are based on context-sensitive 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 in case of shifting contexts or data oscillating near thresholds.

The temporal abstraction methods are integrated and implemented 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.

1. Introduction: the Need for Deriving Temporal Patterns

If one dares to work with monitoring and therapy planning in real-world environments, one faces a host of 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. Additionally, the interpretation context is shifting depending on observed data, and the underlying expectations of the development of parameters are different according to the interpretation context as well as to the degrees of parameters' abnormality.

Theories of data analysis (Avent and Charlton 1990; 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 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 unusable, inappropriate or at least only partially applicable.

To overcome these limitations, qualitative values or patterns are derived and used to improve monitoring and therapy planning. An advantage of using qualitative descriptions is their unified usability in the system model, no matter of what their origin. These derived qualitative values or patterns are used for recommending therapeutic actions as well as for assessing the effectiveness of these actions within a certain period. Several different approaches have been introduced to perform data abstraction (e.g., (Haimowitz, Le, and Kohane 1995; Shahar and Musen 1993) a detailed comparison is given in Section 2). However, dealing with high-frequency data, shifting contexts, and different expectations of the development of parameters require particular temporal abstraction methods to arrive at unified qualitative values or patterns.

We propose context-sensitive and expectation-guided temporal abstraction methods. They incorporate knowledge about data points, data intervals, and expected qualitative

trend patterns to arrive at unified qualitative descriptions of parameters (temporal data abstraction). Our methods are based on context-sensitive 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 in case of shifting contexts or data oscillating near thresholds. Our temporal abstraction methods combine AI techniques with time-series analysis, namely linear regression modeling. The stepwise linear regression model approximates vague medical knowledge, which could be determined only in verbal terms.

Our approach is oriented toward, but not limited to, our application domain: artificial ventilation of newborn infants in intensive care units. The temporal abstraction methods are integrated and implemented in VIE-VENT, an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants (Miksch, et al. 1993). VIE-VENT had been tested and evaluated in real clinical scenarios. The applicability and usefulness of our approach are illustrated by an example of VIE-VENT.

In the first part of this paper we will illustrate why previous methods are not applicable and fail to meet our requirements. The second part will describe the application domain by introducing a sample case and the basic concepts to proceed with our approach. In the third part we will concentrate on the context-sensitive and expectation-guided temporal abstraction methods and illustrate them using our sample case. Finally, we will describe our experiences within a real-clinical setting concluding with strengths and limitations of our approach.

2. Alternative Approaches and their Limitations: the Need for New Data-Abstraction Methods

During the recent years, several different approaches have been introduced to perform temporal abstraction tasks. The systems were implemented mainly for clinical domains. A pioneer work in the area of knowledge-based monitoring and therapy planning systems was the Ventilator Manager (VM, (Fagan, Shortliffe, and Buchanan 1980)), which was designed to manage postsurgical mechanically ventilated patients. VM was developed in the late 1970s as one of a series of experiments studying the effectiveness of the MYCIN formalism. In recent years the most significant and encouraging approaches were the temporal utility package (TUP, (Kohane 1986)), the temporal control structure system (TCS (Russ 1989)), the TOPAZ system (Kahn 1991), the temporal-abstraction module in the M-HTP project (Larizza, Moglia, and Stefanelli 1992), the Guardian project (Hayes-Roth, et al. 1992), the Trendx system (Haimowitz, Le, and Kohane 1995), and RÉSUMÉ (Shahar and Musen 1993; Shahar and Musen 1996). A comprehensive review of temporal-reasoning approaches and useful references are given in (Shahar and Musen 1996). In the fol-

lowing we will concentrate only on the two approaches most closely related to our approach, pointing out their differences and limitations for our purpose.

Haimowitz and Kohane (Haimowitz, Le, and Kohane 1995) have developed the concept of trend templates (Trendx) to represent all available information during an observation process. 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 many domains it is impossible to define such static trajectories of the observed parameters in advance. Depending on the degrees of parameters' abnormalities and on the various contexts, different normal behaviors are expected. These normal expectations vary according to the patient's status in the past. Therefore these thresholds have to be derived dynamically during the observation period. For example, the decreasing of transcutaneous partial pressure of carbon dioxide ($P_{tc}CO_2$) 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 extremely above the target range in the past. However, the same amount of change (4 units) from 54 mmHg to 50 mmHg would be assessed as "normal decrease" during a period where the patient's respiratory status was slightly above the target range.

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 because of 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 some domains. Third, it concentrates on methods to cope with low-frequency observations which cannot easily be adapted for high-frequency data due to their different properties. Fourth, different contexts have to be defined in advance and are not automatically deduced from the input parameters. Fifth, the high level abstraction mechanism (pattern matching based on external and internal knowledge) is superfluous for therapy planning.

Our approach benefits from using all available information 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 methods cover the different degrees of parameters' abnormalities caused by shifting contexts and their corresponding dynamics (e.g., "the higher the degree of a parameter's abnormality the bigger is the amount of positive parameter's change 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. To keep our qualitative descriptions stable we apply smoothing and adjustment methods.

Additionally, we do not predefine absolute, time-dependent expected normal behavior of parameters during the whole observation process (as in (Haimowitz, Le, and Kohane 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. The combination of different parameters' states reflects a particular context. Improving or worsening of these parameters are assumed to be best described as exponential functions. The costs to compare such exponential functions are reduced by stepwise linearization.

3. Application Domain and Basic Concepts

In the following section we will explain our application domain, specify the input and the output of our temporal data-abstraction methods, introduce a sample case, and explain the basic notion of our concepts "context-sensitive" and "expectation-guided".

3.1 Application Domain: Monitoring and Therapy Planning of Artificially Ventilated Newborn Infants in NICUs

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, create a rising information management problem at neonatal ICUs (NICUs). We are particularly interested in the monitoring and therapy-planning tasks of artificially ventilated newborn infants in NICUs. These tasks can be improved by applying derived qualitative values or patterns (temporal data abstraction).

Our temporal abstraction methods are integrated, implemented, and evaluated in VIE-VENT. VIE-VENT is an open-loop knowledge-based monitoring and therapy planning system for artificially ventilated newborn infants (Miksch, et al. 1993; Miksch, et al. 1995). It incorporates alarming, monitoring, and therapy planning tasks within

one system. The data-driven architecture of VIE-VENT consists of five modules: data selection, data validation, temporal 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 NICUs. Its various components are built in analogy to the clinical reasoning process.

3.2 Input and Output

VIE-VENT's input data set can be divided into continuously and discontinuously assessed data. Continuously assessed data (e.g., blood gas measurements, like $P_{tc}O_2$, $P_{tc}CO_2$, S_aO_2 , and ventilator settings, like PIP, F_iO_2) are taken from the output of the data selection module every 10 seconds. Discontinuously assessed 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.

The input of the temporal data-abstraction methods includes a set of time-stamped parameters (the continuously assessed data retrieved every 10 seconds and the discontinuously assessed data at a particular time-stamp) and expected qualitative trend patterns (e.g., "the parameter $P_{tc}CO_2$ is moving one qualitative step towards the target range within 20 to 30 minutes."). The specific context of the observed parameters is automatically deduced from the input parameters, mainly the ventilator settings. The output of the data-abstraction methods is a set of time-point- and interval-based, context-specific, qualitative descriptions. These qualitative descriptions can be a separate abstraction at a particular time-stamp and/or a combination of different time-specific abstractions (a higher level of abstraction, e.g., a combination of different time-stamped qualitative data-point categories or a combination of time-point- and interval-based values called qualitative trend category).

3.3 A Sample Case

Figure 1 shows a sample case of VIE-VENT. In the following sections this sample case will be used to illustrate our temporal data-abstraction methods. The left-hand region shows the blood gas measurements (transcutaneous CO_2 , O_2 , S_aO_2) and their corresponding qualitative temporal abstractions on the top. The actual ventilator settings (first column, e.g., F_iO_2 is 38%), and VIE-VENT's therapeutic recommendations at the current time (second column, e.g., decrease F_iO_2 to 30%) are given below. The upper right-hand region shows two status lines. First, the combination of different time-specific abstractions is labeled by "Status" (e.g., "hyperoxemia" is the combination of the qualitative data-point categories of S_aO_2 and $P_{tc}O_2$). Second, additional warnings are labeled by "Warns" (e.g., "worsening"

means that VIE-VENT detected, that the respiratory system of the neonate is worsening). The right-hand region gives plots of the most important parameters over the last four hours. Scrolling to previous time periods is possible by pushing 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 (pushing the buttons <Plot 2>, <History> and <Trend>, respectively). The therapeutic recommendations are displayed as red vertical lines in the corresponding curve of the ventilator setting.

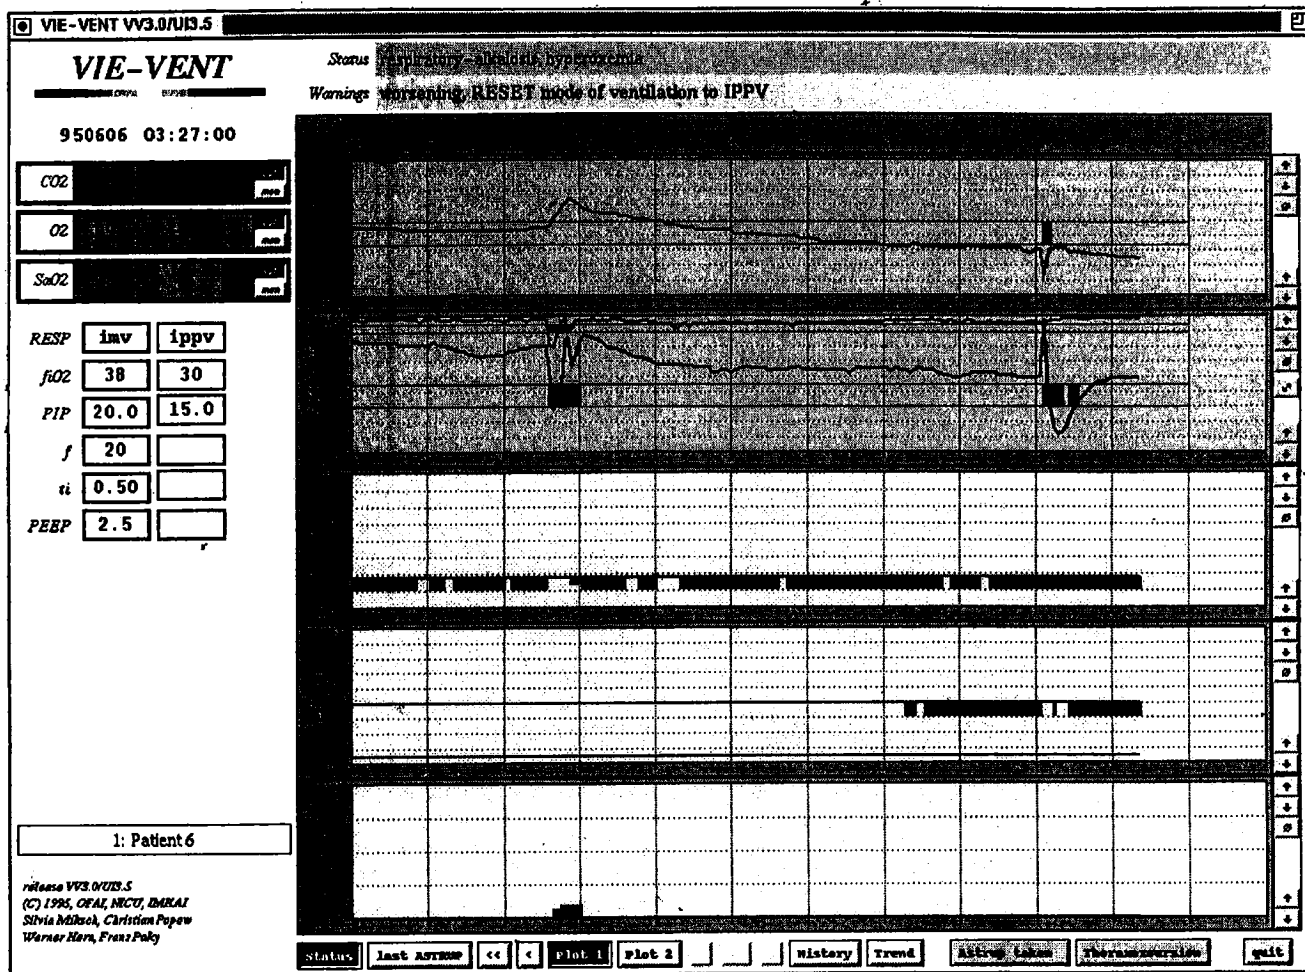


Figure 1: Sample case of VIE-VENT. The left-hand region shows the blood gas measurements, their corresponding qualitative temporal abstractions on the top and the actual and recommended ventilator settings below. The right-hand region gives plots of the most important parameters over the last four hours, namely transcutaneously assessed blood gas measurements and some ventilator settings.

3.4 Meaning of "Context-Sensitive"

The abstraction problem becomes more difficult when the behavior of a system involves interactions among components or interactions with people or with the environment. Under these conditions, correct abstractions become context-sensitive. It is possible to determine *a priori* a set of sensor parameters with their fixed plausible ranges.

However, if the context is shifting, e.g., one component gets in a critical condition or a changing of specific phases or protocols occurs, a capability for dynamic adjustment of threshold values is needed.

The context is automatically deduced from the set of input parameters. For example, we monitor the patient during the whole artificial ventilation process. The ventilation process can be divided into different phases, namely an initial phase, a phase of controlled ventilation (intermittent posi-

tive pressure ventilation, IPPV), a phase of weaning (intermittent mandatory ventilation, IMV), and a phase of returning to spontaneous breathing. All phases characterize a particular context and can be deduced from the current ventilator setting. In Figure 1 the context "imv" is shown in the first row of the ventilator settings labeled by "RESP". The second column gives the current recommendation of VIE-VENT (i.e., change the context to "ippv"). The user interface is designed for physicians. Therefore we used labels which are meaningful for physicians. We defined context-specific transformation schemata of time-stamped data as well as adjustment methods in case of shifting contexts and data oscillating near thresholds.

3.5 Meaning of "Expectation-Guided"

Usually, the temporal abstraction is either exclusively based on the observed input parameters (compare (Shahar and Musen 1993; Shahar and Musen 1996)) or predefined trajectories of observed parameters are used (compare (Haimowitz, Le, and Kohane 1995)). The first neglects available knowledge, in many domains expectations of parameters' courses are obtainable. However, trajectories of observed parameters are often difficult to define in advance. The problem lies in the lack of an appropriate curve-fitting model to predict the development of parameters from actual measurements. Nevertheless, verbal descriptions about expectations of parameters' developments are attainable from domain experts. We improved our temporal data-abstraction process, including *expected qualitative trend descriptions*, which are derived from domain experts.

In the next section we will explain our temporal data-abstraction methods in detail.

4. Temporal Data-Abstraction Methods

The aim of the temporal data-abstraction process is to arrive at unified, context-sensitive qualitative descriptions. The data abstraction is based on time points, time intervals and expected qualitative trend descriptions within a particular context.

Dealing with high-frequency data, shifting contexts, and different expectations of the parameters' development requires particular temporal abstraction methods to arrive at unified qualitative values or patterns. Our temporal data-abstraction process consists of five different methods: (1) transformation of quantitative point data into qualitative values (context-sensitive schemata for data-point transformation), (2) smoothing of data oscillating near thresholds, (3) smoothing of schemata for data-point transformation, (4) context-sensitive adjustment of qualitative values, (5) transformation of interval data (context-sensitive and expectation-guided schemata for trend-curve fitting).

The schemata for data-point transformation transform single observations into qualitative values. To keep the qualitative

values stable in case of shifting contexts or data oscillating near thresholds, we apply different smoothing methods. In critical states of the patient we have to adjust the qualitative values avoiding severe lung damage (context-sensitive adjustment of qualitative values). The schemata for curve fitting represent the dynamically changing knowledge to classify the observed parameters in combination with different expectations of the parameters' courses during time periods. The next sections explain these methods in detail.

4.1 Context-Sensitive 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. The region defines the only common property of the numerical and qualitative values within a particular context and at a specific time-stamp. It is comparable to the "point temporal abstraction" task of Shahar and Musen (Shahar and Musen 1993).

The bases of our transformation of the blood gas measurements are context-sensitive *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 interests 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 4.5). 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 all different contexts (e.g., "imv"). The different contexts require specific predefined target values depending on different attainable goals. Figure 2 shows the schema of transcutaneous partial pressure of carbon dioxide (P_{tCO_2}) during IMV. For example, the transformation of the transcutaneous P_{tCO_2} value of 34 mmHg during IMV results in a qualitative P_{tCO_2} value of *g2* ("substantially below target range") whereas during IPPV it would represent *g1* ("slightly below target range"). 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.

In Figure 1 the temporal abstraction of the blood gas measurements is displayed in the left upper corner. The qualitative data-point categories are expressed using a color chart with different gradation (e.g., deep pink represents values extremely above the target range (s3), lime green represents values extremely below the target range (g3)). The above example of the transcutaneous P_{tCO_2} value of 34 mmHg during IMV is displayed in color chartreuse.

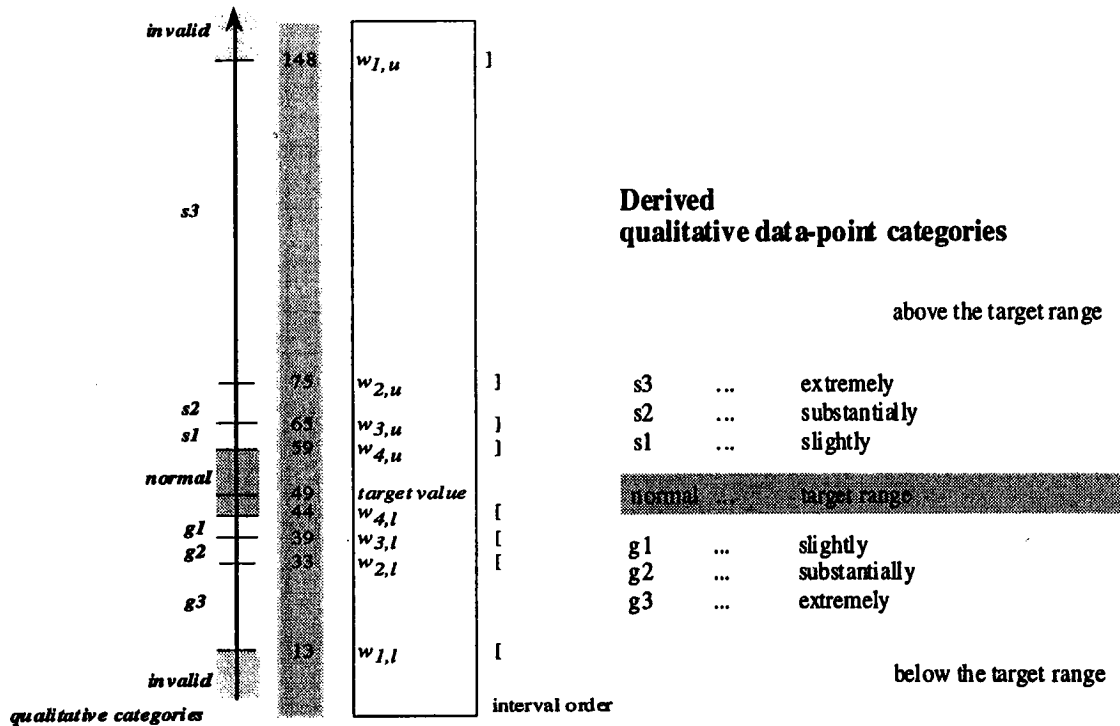


Figure 2: Schema for data-point transformation of $P_{tc}CO_2$ during context IMV. On the left-hand side the abbreviations of the seven derived qualitative data-point categories are used. The labels $w_{i,x}$ indicates the thresholds of the regions of interests. The square brackets [and] show the interval order (e.g.,]75, 148] is a left-side open interval).

4.2 Smoothing of Data Oscillating Near Thresholds

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 method. The key idea is to keep the qualitative categories stable if the quantitative values cross the border to the next qualitative category just minimally for a few moments. Our smoothing method is based on the size of the regions of interests, predefined ϵ regions, and lasting time intervals. Alternative smoothing approaches could use statistical measurements (e.g., interval of confidence) or fuzzy sets to classify the parameter values.

The smoothing method:

Let - in contrast to Figure 2, the second index (upper and lower region) has been eliminated to increase readability - a_t be the actual value at current time t with $a_t \in [w_i, w_{i+1}]$, a_{t-1} be the value on 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
     $((a_t \leq w_i + \epsilon) \text{ or } (a_t \geq w_{i+1} - \epsilon))$ 
then start smoothing
     $\forall a_m, m \in [t, t+x] :$ 

```

```

if  $qual(a_m) \neq qual(a_{t-1})$  and
     $((a_m \leq w_i + \epsilon) \text{ or } (a_m \geq w_{i+1} - \epsilon))$ 
then  $qual(a_m) \leftarrow qual(a_{t-1})$ 
else if  $((a_m > w_i + \epsilon) \text{ or } (a_m < w_{i+1} - \epsilon))$ 
then stop smoothing

```

```

with if  $|w_{i+1} - w_i| > 3$  then  $\epsilon = 2$ 
    else  $\epsilon = 1$ 

```

and $[t, t+x]$ be the lasting time interval

The smoothing method 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 lasting 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 3 gives an example of our smoothing method. At time point t the smoothing method 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 the distance between a_{t-1} and a_{t+6} is greater ϵ . In this example, the predefined time period of 5 minutes has not been exhausted.

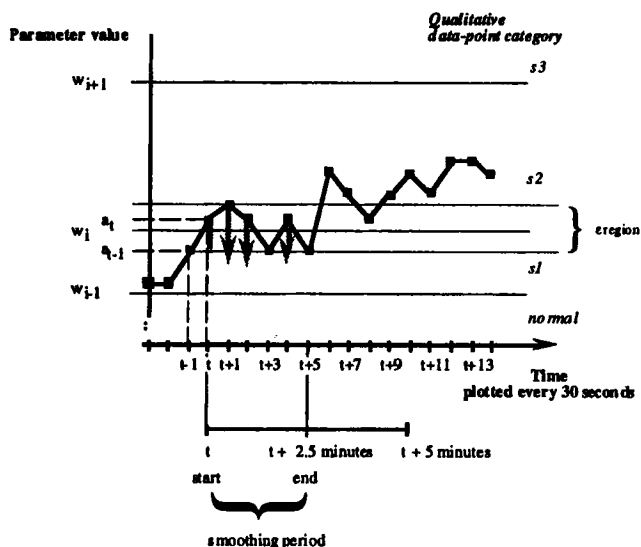


Figure 3: Example of data smoothing.

4.3 Smoothing of Data-Point Transformation Schemata

The schemata for data-point transformation are defined for all contexts (i.e., modes of ventilation: IMV, IPPV) representing different target values. Changing context 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 the weaning process.

4.4 Context-Sensitive Adjustment of Qualitative Values

For extremely critical or life-threatening situations, 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 rule represents this kind of knowledge.

```

if (30 < PIP ≤ 35) and
      (PtCCO2 is "extremely below target range")
then (PtCCO2 is changed to "substantially below target range")
  
```

4.5 Transformation of Interval Data (Context-Sensitive and Expectation-Guided Schema for Trend-Curve Fitting)

Similar 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 for how a blood gas value has to change over time to reach the target range in a physiologically proper way are expressed in verbal terms. For example, "the parameter P_{tC}CO₂ is moving one qualitative step towards the target range within 20 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 4.1. Combining these qualitative data-point categories with the expected qualitative trend descriptions we reach 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 classifies a faster positive change of this parameter as normal. The different sizes of the data-point categories express this circumstance. The corresponding dynamically derived trends depending on the expected qualitative trend descriptions represent different dynamic changes.

Based on physiological criteria, four kinds of trends of our 10-second 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 (Miksich, et al. 1994)).

The transformation of interval data into qualitative values is the last step of the temporal 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 used trend 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.

4.5.1. The Guiding Principle

The guiding principle of our approach is illustrated in Figure 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 for 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 fitted by exponential functions. An appropriate approach classifying trend data is to transform the curve (borders of the dark gray area) shown in Figure 4 into an exponential function and to compare it with the actual growth rate. To classify the trend data, we used a dynamic comparison algorithm which performs a stepwise linearization of the expected exponential function to overcome complexity (compare Section 4.5.2). For example, if a $P_{tc}CO_2$ data point during the context "IMV" 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.

4.5.2 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 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 (at

thresholds horizontal dotted lines are drawn). Their corresponding qualitative categories are listed on the right-hand side. Based on the guiding principle depicted in Figure 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).

The growth rates are calculated and classified for all kinds of trends (very-short-, short-, medium-, and long-term). To increase readability, we show only the principal method and not the results for the four kinds of trends. The algorithm works the same way 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 for the growth rate k_1 and k_2 depending on the relative position of 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 on the qualitative region where the previous data point (a_{t-1}) belongs. In addition to k_1 and k_2 we use an ϵ range around zero 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. The target range of a parameter divides the qualitative regions into an upper part (A1, A2, A3, ZA, C) and a lower part (B1, B2, B3, ZB, D) as explained in Figure 4. The classification process results in instantiations of qualitative trend descriptions for each blood gas measurement, for each kind of trend, and for each activated context.

In Figure 1 the qualitative trend categories are visualized by colored arrows next to the qualitative data-point categories. The four arrows show the directions of the very-short, short, medium, and long-term trends. For example, all qualitative trend categories of $P_{tc}CO_2$ during the context "IMV" are derived as "D" (their directions are down-going and the color is deep-pink). This expresses a dangerous decrease of the measurement. Consequently, our therapy planning module recommends a therapeutic action to decrease PIP (compare fourth plot on the right-hand side in Figure 1). The qualitative trend categories for the short-term trend (second arrow) of $P_{tc}O_2$ and S_aO_2 are derived as "ZA" (zero change) and "A3" (slow decrease), respectively. For S_aO_2 we see a short-term trend of slow decrease, but a zero change during the last 30 minutes (third arrow) and a dangerous increase during the last 3 hours (red fourth arrow pointing upwards). This knowledge is used in our therapy planning module to recommend therapeutic actions. In this case a therapeutic action to decrease F_iO_2 is recommended (compare third plot on the right-hand side in Figure 1).



- Derived qualitative trend categories:**
- 1) A1 ... decrease too fast
 - 2) A2 ... normal decrease
 - 3) A3 ... decrease too slow
 - 4) ZA ... zero change
 - 5) C ... dangerous increase
 - 6) B1 ... increase too fast
 - 7) B2 ... normal increase
 - 8) B3 ... increase too slow
 - 9) ZB ... zero change
 - 10) D ... dangerous decrease

Figure 4: Schema for trend-curve fitting of $P_{tc}CO_2$. 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.




Figure 5: Example of schema for trend-curve fitting of $P_{tc}CO_2$.

5. Applicability and Practical Usefulness

We have tested the applicability of our approach both on generated data sets and on real data. The generated data sets were used to simulate extreme cases. The results obtained demonstrated the robustness of VIE-VENT. Real data were obtained from a NICU using on-line data acquisition. We collected sequences of 16-28 hours of continuous recording of transcutaneous blood gas measurements and pulseoxime-

try. Discontinuously assessed data were taken from the computer-based patient records. The evaluation of these cases demonstrated the applicability of our approach in the clinical setting.

The usefulness of the qualitative categories and their visualizations have been manifested in different ways. First, they support the physicians to get a closer insight into their medical reasoning process. This has eased the fine-tuning of our therapy planning component. Second, the qualitative trend categories improved our data validations component. Third, applying the qualitative trend categories for formulating and assessing therapeutic actions resulted in a graceful weaning process avoiding too abrupt changes of therapeutic recommendations. In Figure 1 the therapeutic recommendations are displayed as red vertical impulses in the corresponding plot of the ventilator setting. The therapeutic recommendations show a very consistent and reasonable picture, except in cases where the measurements were set invalid (gray areas between the two horizontal lines in the two upper plots in Figure 1).

During our evaluation phase we discovered also limitations of our temporal data-abstraction methods. First, information about the frequency of temporal abstractions in the past (e.g., "three episodes of hyperoxemia during the last 3 hours occurred") would be very useful for future reasoning processes. Second, dealing with real data during longer time periods has to take into account that more recently observed data are more important for the reasoning process than data observed in older time periods. Therefore, the data-abstraction methods have to include a memory which weights the time-ordered data.

6. Conclusion

We demonstrated very powerful temporal data-abstraction methods, which combine all available information to perform a context-sensitive and expectation-guided temporal abstraction process. Designing our abstraction, we concentrate on knowledge-based monitoring and therapy planning in real clinical environments. Dealing with high-frequency data, shifting contexts, and different expectations of the development of parameters requires particular temporal abstraction methods to arrive at unified qualitative values or patterns. Our temporal data-abstraction methods incorporate knowledge about data points, data intervals, and expected qualitative trend patterns. Additionally, 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 is involved. Therefore theories of data analysis are only partially applicable. We overcome these limitations applying qualitative statements (called *expected qualitative trend descriptions*), which are obtainable from domain experts. These qualitative statements are approximated using linear regression models. To keep the qualitative descriptions stable in case of shifting contexts or data oscillating near thresholds we apply smoothing and adjustment methods.

Integrating the temporal abstraction methods in VIE-VENT results in easily comprehensible and transparent definitions of the data-interpretation, therapy-planning, and data validation modules. The data interpretation can be performed on different levels using data-point and data-interval (trend) abstractions as well as a combination of different abstraction categories. The derived qualitative values and patterns are used for recommending therapeutic actions as well as for assessing the effectiveness of these actions within a certain period. Additionally, the data validation could be extended using the derived qualitative values and patterns (applying an assessment procedure based on qualitative descriptions). The clinical experiences show that the enhancement of our temporal data-abstraction methods has improved our therapy planning component remarkably. They guarantee a graceful weaning process, avoiding too abrupt changes of parameters.

Acknowledgment. This phase of the project was supported by the "Jubiläumsfonds der Oesterreichischen Nationalbank", Vienna, Austria, project number 4666. Currently, future research is supported 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 (OFAI) by the Austrian Federal Ministry of Science and Research, Vienna.

References

Allen, J.F. 1991. Time and Time Again: The Many Ways to Represent Time. *International Journal of Intelligent Systems*, 6:341-55.

Avent, R.K. and Charlton, J.D. 1990. A Critical Review of Trend-Detection Methodologies for Biomedical

Monitoring Systems. *Critical Reviews in Biomedical Engineering*, 16(6):621-59.

Dean, T.L. and McDermott, D.V. 1987. Temporal Data Base Management. *Artificial Intelligence*, 32(1):1-55.

Fagan, L.M., Shortliffe, E.H. and Buchanan, B.G. 1980. Computer-Based Medical Decision Making: from MYCIN to VM. *Automedica*, 3:97-106.

Haimowitz, I.J., Le, P.P., and Kohane, I.S. 1995. Clinical Monitoring Using Regression-Based Trend Templates. *Artificial Intelligence in Medicine*, 7(6):473-96.

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):165-85.

Kahn, M.G. 1991. Combining Physiologic Models and Symbolic Methods to Interpret Time-Varying Patient Data. *Methods of Information in Medicine*, 30(3):167-78.

Kay, S.M. 1993. *Fundamentals of Statistical Signal Processing* New Jersey:PTR Prentice Hall:Englewood.

Kohane, I.S. 1986. Medical Reasoning in Medical Expert Systems. In Salamon, R., et al., (eds.), *Proceedings of the Fifth Conference on Medical Informatics (MEDINFO-86)*, 170-4. North-Holland:Amsterdam.

Larizza, C., Moglia, A., and Stefanelli, M. 1992. M-HTP: A System for Monitoring Heart Transplant Patients. *Artificial Intelligence in Medicine*, 4(2):111-26.

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., et al. (eds.), *Proceedings of the Artificial Intelligence in Medicine Europe (AIME-93)*, 218-29. IOS Press:Amsterdam.

Miksch, S., Horn, W., Popow, C., and Paky, F. 1994. 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)*, Amsterdam, 48-52. Wiley:Chichester, UK.

Miksch, S., Horn, W., Popow, C., and Paky, F. 1995. Therapy Planning Using Qualitative Trend Descriptions. In Barahona, P., et al. (eds.), *Proceedings of the Artificial Intelligence in Medicine, 5th European Conference on Artificial Intelligence in Medicine Europe (AIME-95)*, Pavia, Italy, 197-208. Springer:Berlin.

Russ, T.A. 1989. Using Hindsight in Medical Decision Making. In Kingsland, L.C. (ed.) *Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-89)*, 38-44. IEEE Computer Society Press:Washington D.C.

Shahar, Y. and Musen, M.A. 1993. RÉSUMÉ: A Temporal-Abstraction System for Patient Monitoring. *Computers and Biomedical Research*, 26(3):255-73.

Shahar, Y. and Musen, M.A. 1996. Knowledge-Based Temporal Abstraction in Clinical Domains. *Artificial Intelligence in Medicine, Special Issue Temporal Reasoning in Medicine*, forthcoming.

Analyzing Discontinuities in Physical System Models

Pieter J. Mosterman and Gautam Biswas

Center for Intelligent Systems

Box 1679, Sta B

Vanderbilt University

Nashville, TN 37235.

pjm,biswas@vuse.vanderbilt.edu

Abstract

Physical systems are by nature continuous. Perceived discontinuities in system models are in reality nonlinear behaviors that are linearized to reduce computational complexity and prevent stiffness in system analysis and behavior generation. Discontinuities arise primarily from component parameter simplification and time-scale abstraction. Discontinuities in models are handled by introducing idealized switching elements controlled by finite state automata to bond graph models. In this paper, we make a systematic study of the nature and effects of discontinuities in physical system models, and present an algorithm that generates consistent and correct behaviors from hybrid models. A primary contribution of this paper is the characterization of discontinuous changes, and a systematic method for validating system models and behavior generation algorithms.

Introduction

Physical systems are by nature continuous, but often exhibit nonlinearities that make behavior generation complex and hard to analyze especially with qualitative reasoning schemes. This complexity is often reduced by linearizing model constraints so that component models (e.g., transistors, oscillators) exhibit *multiple piecewise linear* behavior or by abstracting time so that component models (e.g., switches, valves) exhibit *abrupt discontinuous* behavior changes[2, 11]. In either case, the physical components are modeled to operate in *multiple modes*, and system models exhibit mixed continuous/discontinuous behaviors.

Typically, simulation methods for generating physical system behavior, whether quantitative (e.g., [12]) or qualitative (e.g., [5, 8]) impose the *conservation of energy* principle and *continuity constraints* to ensure that generated behaviors are meaningful. Energy distribution among the elements of the system defines system *state*, and energy distribution over time reflects the history of system behavior.

Bond graphs[12] provide a general methodology for modeling physical systems in a domain independent way. Its primary elements are two energy storage elements or *buffers*, called capacitors and inductors,¹ and a third element, the resistor, which *dissipates* energy to the environment[10, 12]. Two other elements, sources of flow and effort, allow transfer of energy between the system and its environment. In addition, the modeling framework provides *idealized junctions* that connect sets of elements and allow lossless transfer of energy between them. Two other *specialized junctions*, transformers and gyrators, allow conversion of energy from one form to another. The use of the bond graph modeling language in building *compositional models* of complex systems has been described elsewhere[1].

State changes in a system are caused by energy exchange among its components, which is expressed as *power*, the time derivative, or *flow*, of energy. Power is defined as a product of two conjugate *power variables*: *effort* and *flow*. Given conservation of energy holds for a system, the time integral relation between energy variables and power variables implies *continuity* of power, and, therefore, effort and flow. As discussed earlier, discontinuities in behavior are attributed to time scale abstraction and parameter abstraction.

Consider an example of an ideal non-elastic collision in Fig. 1. A bullet of mass m_1 and velocity v_1 hits an unattached piece of wood of mass m_2 , initially at rest. The collision causes the bullet to get implanted in the wood, forming one body of mass $m_1 + m_2$ moving with velocity say v_2 . The initial situation is represented as two bond graph fragments. Each fragment has an inertial element (mass) connected by a 1-junction to a zero effort source (S_e) which indicates that there are no forces acting on either the bullet or the wood. Because there are no sources, the momentum after collision is conserved, so $m_1 v_1 = (m_1 + m_2) v_2$, i.e., $v_2 = \frac{m_1}{m_1 + m_2} v_1$.

¹Capacitors store energy in the form of *generalized displacement* and inductors store energy as *generalized momentum*