ReMINE

High performances prediction, detection and monitoring platform for
patient safety risk management

FP7 Contract: 216134

Third Release of the Adverse Risk Management Support System

Deliverable 4.8 incorporating D4.7

Document Information

Document Name:	REMINE_D4.8_TUW_WP4_v2.3_101025
Deliverable:	D4.8 Third Release of the Adverse Risk Management Support System
Revision:	V2.3
Revision Date:	25/10/2010
Author	Theresia Gschwandtner, Katharina Kaiser, Silvia Miksch, Andreas Seyfang
Security	Public

Diffusion list

Consortium

Approvals

	Name	Beneficiary	Date	Visa
Project Coordinator	Michele CARENINI	NOEMALIFE		
Technical Manager	Fernando GUMMA	NOEMALIFE		

Documents history

Revision	Date	Modification	Author
V0.1	09/06/2009	Creation	Theresia Gschwandtner
V0.2	10/06/2009	Deliverable Refinement	Katharina Kaiser
V0.3	12/06/2009	Deliverable Refinement	Theresia Gschwandtner
V0.4	13/06/2009	Deliverable Refinement	Andreas Seyfang
V0.5	13/06/2009	Revision	Theresia Gschwandtner
V0.6	21/06/2009	Revision	Theresia Gschwandtner
V1.0	16/12/2009	Creation Second Release	Theresia Gschwandtner
V1.1	18/12/2009	Deliverable Refinement	Silvia Miksch
V1.2	20/12/2009	Deliverable Refinement	Andreas Seyfang

V1.3	21/12/2009	Deliverable Refinement	Katharina Kaiser
V1.4	21/12/2009	Revision	Theresia Gschwandtner
V1.5	30/12/2009	Revision	Theresia Gschwandtner
V1.6	15/09/2010	Creation Third Release	Theresia Gschwandtner
V1.7	21/09/2010	Deliverable Refinement	Silvia Miksch
V1.8	23/09/2010	Deliverable Refinement	Katharina Kaiser
V1.9	27/09/2010	Deliverable Refinement	Andreas Seyfang
V2.0	28/09/2010	Revision	Theresia Gschwandtner
V2.1	20/10/2010	Integration of Feedback from ICCS	Theresia Gschwandtner
V2.2	21/10/2010	Revision	Andreas Seyfang
V2.3	25/10/2010	Revision	Theresia Gschwandtner

Table of Contents

1. Executive summary

In recent years sophisticated approaches have been developed in order to permit computerized processing and visualizing of clinical practice guidelines. Two major challenges of the visualization of clinical guidelines are (1) to properly support the time-oriented process of treatment plans, and (2) to properly represent the multidimensional information spaces pertaining to a patient's condition. For both tasks a variety of Information Visualization techniques have been developed; however, visualizing the various interdependencies between the execution of a treatment plan and the patient's condition is still a largely unresolved task.

This deliverable deals with the analysis of requirements to visualise treatment processes, patient data, and combinations of both, with a focus on satisfying the requirements of the ReMINE project. It contains furthermore a comprehensive survey of existing visualization approaches dealing with these aspects as well as a critical review of these approaches. Subsequently, we present the interactive visualization module supplied by TUW, especially tailored towards the needs of RAPS assessment and prevention together with an outline of the interactive mechanisms. Moreover, we give a detailed discussion on design decisions.

This deliverable presents the third version of the Adverse Risk Management Support System. The first version was presented in D4.7 which is incorporated in this document.

2. Introduction

The two main objectives of the ReMINE project are:

(1) to develop a new technological platform that is able to perform semi-automated RAPS (Risks Against Patient Safety) management, and

(2) to propose organisational changes with considerable added value in relevant environments.

The first objective is being achieved by developing the ReMINE system, a RAPS identification and analysis system for the acquisition and mining of relevant multimedia data present in hospitals. This system will then be used to predict, detect and monitor RAPS related events in the collaborating facilities.

Tasks to achieve the second objective include developing adequate clinical risk management processes, establishing a more active role for RAPS managers, and identifying new ways for interactions amongst different health care professionals in a local health care system to solve RAPS issues.

In both tasks an important means is to visualize data and information in order to support the medical staff. Thereby, methods are not only restricted to user interfaces for risk managers, but we are also introducing advanced visualization methods.

Especially in the discipline of medicine the amount of information grows constantly, thus, the application of Information Visualization (InfoVis) techniques to represent the available information in a comprehensible way is becoming imperative. Card et al. constituted Information Visualization as ”The use of computer-supported, interactive, visual representations of abstract data to amplify cognition” [Card et al., 1999]. Providing interactive visual representations, Information Visualization aims to support intelligibility and optimal use of enormous amounts of complex and often interrelated data. “InfoVis is concerned with the development of interactive visual representations of abstract, multidimensional data, information, and knowledge to help users gain a deeper understanding of the contents of a domain by revealing, for example, new insights, previously unknown facts and relationships, or providing explanations for complex situations" [Aigner et al., 2008]. Furthermore, it allows the user to explore data at different stages of granularity to avoid overload, e.g., providing means for both a compact overview and a deep exploration of the "information space".

In the medical domain visualization techniques have been used for many years quite frequently, e.g., 3D volume visualization tasks, X-ray, or computer tomography visualizations. These kinds of applications, i.e., visualizing objects of the physical world, are associated with the discipline of Scientific Visualization. In contrast to this, Information Visualization focuses on the visualization of abstract data, i.e., data for which a natural mapping to physical objects may not exist. Containing mainly abstract data, e.g., patient data, or treatment processes, the visualization of computerized guidelines is a challenge for Information Visualization researchers, which has not been targeted until quite recently [Aigner et al., 2008].

2.1. Role within the Project

In the context of preventing Risks Against Patient Safety using the ReMINE platform, there are several occasions on which information about the patient or patient groups and processes relating to them need to be displayed.

The Guideline Execution Engine produces treatment recommendations based on real or synthesized patient data. These recommendations need to be compared to the treatment steps actually performed. Since this matching is a complex task, a significant part of it will be performed by a human user (the risk manager or a physician supervising the treatment). To assist him in this task, versatile and easily comprehensive representations of complex user data together with process information (showing the treatments) is required.

The ReMINE Metadatabase may allow the risk manager to browse user data and explore the information about a specific patient. After selecting a patient interactively, he might wish to use Information Visualization techniques to browse the data recorded for this patient.

Machine Learning may produce sets of patients which exhibit specific attributes. These attributes will in our case be linked to risk factors or indicators of malpractice or more generally ratings of quality of care. Since most Machine Learning algorithms use some form of matching function, the result of this function produces a numeric value showing the strength of the relation. This could be used to sort the matched patients in the list shown to the risk manager.

The work described in this project must be seen as a concrete vision for the future, rather than an integrated part of the software system, for the several reasons. First of all, electronic patient records with rich sets of integrated quantitative temporal data are not a reality in today's care practice. They are the near future, but not the present or past. Second, pilot site have not been selected based on the richness of the data available there, or the role which protocol-based care plays in their daily routine. Third, most of today's protocols are not data-rich - they leave most data interpretation to the user, which means that the data entered to the system is simple (Boolean or a qualitative value). This is particularly true for low-frequency domains like general disease management, and such was selected in this project. In high-frequency domains, like intensive care units, treatment protocols are generally data-rich, but data integration is much more complex than in the applications chosen in this project, because timing of data transfer is more critical and there is less standardization.

At the same time, scientific literature clearly shows that there is a huge demand for advance visualizations in the near future [Card et al., 1999][Fekete et al., 2008][Keim, 2002].

As a consequence of the status of this work, this deliverable contains two chapters exemplifying the usage of our ideas. Section 8 describes the features in detail, using data from a data-rich domain, the artificial ventilation of neonates. Section 9 describes the visualization of data generated based on the treatment protocols modeled in the ReMINE project.

2.2. Focus of the Work

Within the requirements described above, we face different levels of sophistication and different challenges.

For simple cases, i.e., for one or few parameters, visualization of patient groups is simple and already implemented in Weka, which can be used within the Machine Learning task.

For complex data (i.e., highly structured data, and/or temporal data), showing data from more than one patient inevitably leads to very complex representations. To avoid them, there are two alternatives. a) Selecting a single parameter as the current focus of interest and displaying that for the whole patient group, which turns the problem into a simple one already solved by Weka. b) Selecting one patient from the list. This turns the problem into showing complex data for a single patient.

For complex, time-annotated data pertaining to a single patient (or a small number of patients), we believe that complementing the existing approaches by new research is both dearly needed and fruitful. Doing so, we built on existing expertise in the working group at TUW as well as on external scientific publications.

2.3. Relation to other Deliverables

This deliverable is the third part of the series
D4.7 First Release of the Adverse Risk Management Support System due M18.
ID Second Release of the Adverse Risk Management Support System due M24.
D4.8 Third Release of the Adverse Risk Management Support System due M33.

Another part related to the Risk Management System is the Guideline Execution Engine. It is described in D4.6 "Documentation of Prototypical Implementation of the Guideline Execution Engine", due M20 and M26.

Therefore, the content of the deliverable you are reading is reduced to the description of the advanced visualization techniques developed for the purpose of ReMINE.

2.4. Deliverable Structure

This document is structured as follows. In Section 3 we start with an in-depth analysis of the different kinds of data and information we have to deal with when visualizing the execution of computerized guidelines. We outline visualization-relevant characteristics of treatment plans in Section 3.1 and take a close look on characteristics of information relevant when visualizing a patient’s condition in Section 3.2.

Subsequently, in Section 4 we give a list of visualization challenges based on the results of the data analysis. In Section 5 we give a comprehensive overview of existing visualization approaches; we start with different approaches to visualize treatment plans in Section 5.1, whereupon, we outline selected visualization approaches dealing with the representation of patient data in Section 5.2, and finally, we give a description of three approaches which visualize both, treatment plans and patient data in combination in Section 5.3. In Section 5.4 we give a short description of the Microsoft Health Common User Interface and its applicability and in Section 5.5 we outline relevant techniques to visualize hierarchical data.

In Section 6 we illustrate how well the introduced visualization techniques meet the given requirements and outline some shortcomings.

In Section 7 we sum up the visualization requirements with respect to the findings of analysis of data characteristics and existing visualization approaches.
In Section 8 we present our interactive visualization module and we discuss all design decisions in detail in Section 8.2. Moreover, we give an outline of interactive means which allow for the exploration of data in Section 8.3 . Whereupon, we show example visualizations of ReMINE protocols in Section 9. Subsequently, we outline the results of the usability evaluation we conducted to ensure the quality of the visualization module in Section 10. We describe the integration of the visualization module with the ReMINE project in Section 11. Finally, we conclude this deliverable in Section 12.

3. Analysis of the Data

When visualizing the execution of clinical practice guidelines (CPGs), there are two interconnected sets of data and information to be considered:

1. The progress of treatment plans and
2. The course of the patient.

On the one hand it is crucial to visualize the currently applied treatment plan, what therapy or medication is applied at the moment, which treatments have been applied so far, etc. Besides the progress of the treatment, the internal structure of these plans is to be communicated by the visualization. Treatment plans usually consist of different sub-plans, which again may contain sub-plans, etc. These plans and sub-plans have to be executed in a strictly defined order which should be visualized as well.
It is understood that the applicable treatment plans very much depend on the patient's condition. Thus, the course of the patient in connection with the applied treatment plans has to be represented in an intelligible way. Due to the use of modern information technologies in clinical care, the amount of information grows constantly. Containing heterogeneous information spaces, i.e., quantitative and qualitative data, high-frequency and low-frequency data, a mixture of text, images, and numerical values, etc., patient data conceptually can be seen as multidimensional information spaces. Consequently, sophisticated InfoVis techniques have to be applied to provide an information rich overview of these enormous amounts of heterogeneous data and information generated for a single patient.

3.1. Visualization of Treatment Plans

Using the Asbru guideline representation language [Miksch et al., 1997a][Seyfang et al., 2002] (used in the Asgaard Project [Shahar et al., 1998]) to express clinical protocols or CPGs, we have to analyze the type and structure of Asbru plans in order to find a suitable visual representation, as treatment plans formulated in Asbru yield some characteristics relevant for visualization:

- Hierarchical decomposition;
- Logical sequences and execution order; and
- Complex plan characteristics.

3.1.1. Hierarchical Decomposition

An Asbru plan consists of different element types: *preferences, intentions, conditions, effects, and a plan body*.
The **plan body** is decomposed into sub-plans, which are performed if given preconditions of the plan hold. Sub-plans must be performed according to the plan's type: in sequence, in any order, in parallel, unordered, or periodically (see Section 3.1.2). Only one type of plan is allowed in a single plan body. Non-decomposable plans are called user-performed plans or actions. All the sub-plans consist of the same components as the plan, namely: preferences, intentions, conditions, effects, and the plan body itself [Kosara & Miksch, 2001]. Thus, a plan and its sub-plans form a hierarchical structure, which is relevant for execution.

3.1.2. Logical Sequences and Execution Order

According to the execution sequence of the sub-plans constituting a plan, we distinguish between several types of plans: *sequential plans, parallel plans, any-order plans, unordered plans, and cyclical plans*.

- All sub-plans of a **sequential plan** are to be executed one after the other in a predefined order.
- All sub-plans of a **parallel plan** get initialized at the beginning and are processed in parallel.
- All sub-plans of an **any-order plan** are processed one after the other but the order is arbitrary.
- All sub-plans of an **unordered plan** can be executed in an arbitrary way – in sequence or parallel or some in sequence and others in parallel. The order is arbitrary.
- A **cyclical plan** contains a sub-plan that can be repeated. The time interval over which the plan is repeated, the start time, the end time, and the duration of the particular plan instance in each cycle are determined.

3.1.3. Complex Plan Characteristics

Preferences define the applicability of a plan, i.e., it serves as a select criteria how well the plan fits the patient’s condition (e.g., exact fit, roughly fit), as well as express the behavior of the plan, for instance, the kind of strategy – aggressive or normal.
Intentions are high-level goals at various levels of the plan. They are temporal patterns of executing-agent actions and external-world states that should be maintained, achieved, or avoided.
We define four categories of intentions [Miksch et al., 1997a]:

1. **Intermediate state**: the state(s) that should be maintained, achieved, or avoided during the applicability of the plan (e.g., weight gain levels are slightly low to slightly high);
2. **Intermediate action**: the action(s) that should take place during the execution of the plan (e.g., monitor blood glucose once a day);
3. **Overall state pattern**: the overall pattern of states that should hold after finishing the plan (e.g., patient had less than one high glucose value per week);
4. **Overall action pattern**: the overall pattern of actions that should hold after finishing the plan (e.g., patient had visited dietitian regularly for at least three months).

This information is important not only for selecting the right plan but also for critiquing treatment plans as part of the ever ongoing process of improving the treatment. This makes intentions one of the key parts of Asbru [Miksch et al., 1997b][Kosara & Miksch, 2001].

Conditions need to hold at particular plan steps in order for a plan to be started, suspended, reactivated, aborted, or completed (see Figure 1), for instance, preconditions, suspension conditions, or abort conditions. Two different kinds of preconditions exist, that must be true in order for a plan to be started: filter-preconditions cannot be achieved through treatment (e.g., subject is female); setup-preconditions can. After a plan has been started, it can be suspended until either the restart-condition is true (whereupon it is continued at the point where it was interrupted) or it has to be aborted. If a plan is aborted, it has failed to reach its goals. If a plan completes (the completion conditions hold) it has reached its goals and the next plan in the sequence is to be executed [Kosara & Miksch, 2001].

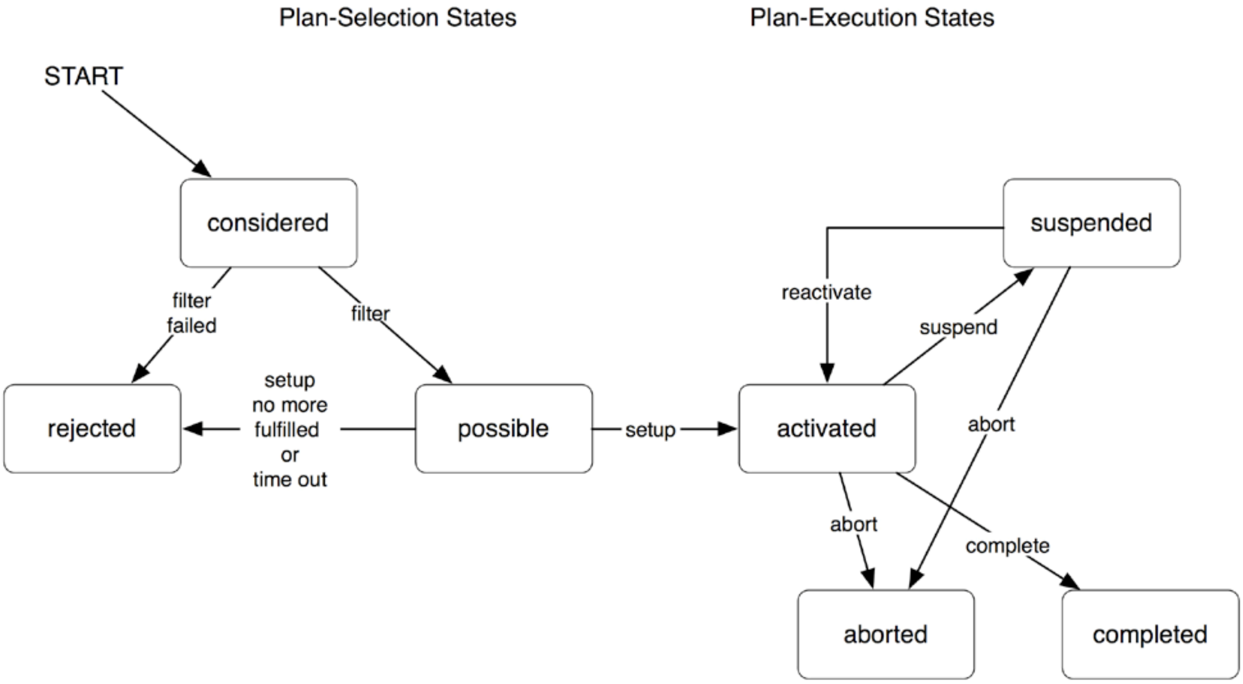

Figure 1: Different states of Asbru plans.

Effects describe the relationship between plan arguments and measurable parameters by means of mathematical functions. A probability of occurrence is also given [Kosara & Miksch, 2001].

3.2. Visualization of a Patient's Condition

Due to the use of modern information technologies in clinical care, the amount of information grows constantly. Thus, the application of Information Visualization techniques to deal with the available information in a comprehensible way is imperative.

Being composed of many heterogeneous information spaces, i.e., including quantitative and qualitative data, high-frequency and low-frequency data, a mixture of text, images, and numerical values, etc, patient data can be seen as multidimensional information spaces. In particular, because of this heterogeneity and because it contains mainly abstract data, the visualization of a comprehensive picture of a patient's condition and especially its evolution over time is a great challenge for Information Visualization researchers [Aigner et al., 2008].

4. Visualization Challenges

A feasible visualization of the course of a patient in connection with the execution of clinical protocols or guidelines implies the communication of all important data and information. With respect to the characteristics of the above mentioned data and information, we face the following challenges:

1. Visualizing the hierarchical structure of treatment plans and sub-plans
2. Visualizing the execution sequence of treatment plans
3. Visualizing relevant characteristics of treatment plans (e.g., plan states, conditions, and effects)
4. Drawing a comprehensive picture of a patient's condition
5. Finding a compact but comprehensive visualization method for the patient's condition and its evolution over time.
6. Visualizing applied treatment plans to make interrelations between them transparent.

5. Existing Approaches

In the following subsections we outline existing approaches focusing on the visualization of treatment plans, patient data, or both patient data in combination with treatment plans.

5.1. Visualization of Treatment Plans

The complex nature of clinical practice guidelines demands for a plain and compact visualization of the underlying information. Dealing with logical sequences, hierarchical data, as well as time-oriented data, the visualization of clinical treatment plans relates to several specific fields of Information Visualization.

In recent years different frameworks to implement computerized guidelines have been developed, such as Asbru, EON, GLIF, Guide, Prodigy, and PROforma. For each framework various visualization methods have been developed in order to communicate the complex underlying information.

5.1.1. Clinical Algorithm Maps

Clinical algorithm maps are also called *flowchart algorithms* and are the most widely used representation of clinical guidelines or treatment plans. They were intended for paper use and have never been enriched by computer support (e.g. navigation, or annotation possibilities). The *Committee on Standardization of Clinical Algorithms* of the *Society for Medical Decision Making* has proposed a standard for clinical algorithm representation [Society for Medical Decision Making, 1992]:

"However, since algorithmic logic is wired implicitly into a protocol, it is difficult to learn an algorithm from a protocol. By contrast, flow-chart algorithms or clinical algorithm maps, are uniquely suited for explicitly communicating conditional logic and have therefore become the main format for representing a clinical algorithm clearly and succinctly."

This standard includes a small number of different symbols as well as some rules on how to use them. One extension to standard flowcharts are annotations including further details: citations to supporting literature, clarifications for the rationale of decisions, etc. The most important advantage of clinical algorithm maps is that they are well known among physicians and thus, requiring minimal learning effort. One drawback is the immense space consumption when representing more complex situations. This leads to huge diagrams where overview is lost easily. Moreover, clinical algorithm maps do neither support temporal information besides “before” and “after”, nor concurrent tasks and complex conditions.

5.1.2. Protégé

Protégé [Gennari et al., 2002] is a free, open source ontology editor and knowledge-base framework. The extensible development environment allows for ontology development and knowledge acquisition in order to facilitate the authoring of clinical guidelines in various guideline representation languages (e.g., EON, GLIF, Prodigy, and PROforma). The graphical user interface illustrates the clinical algorithm in a way similar to flowcharts by using different shapes for plans, decisions, actions, enquiries, and root tasks which are connected by arrows (see Figure 2).

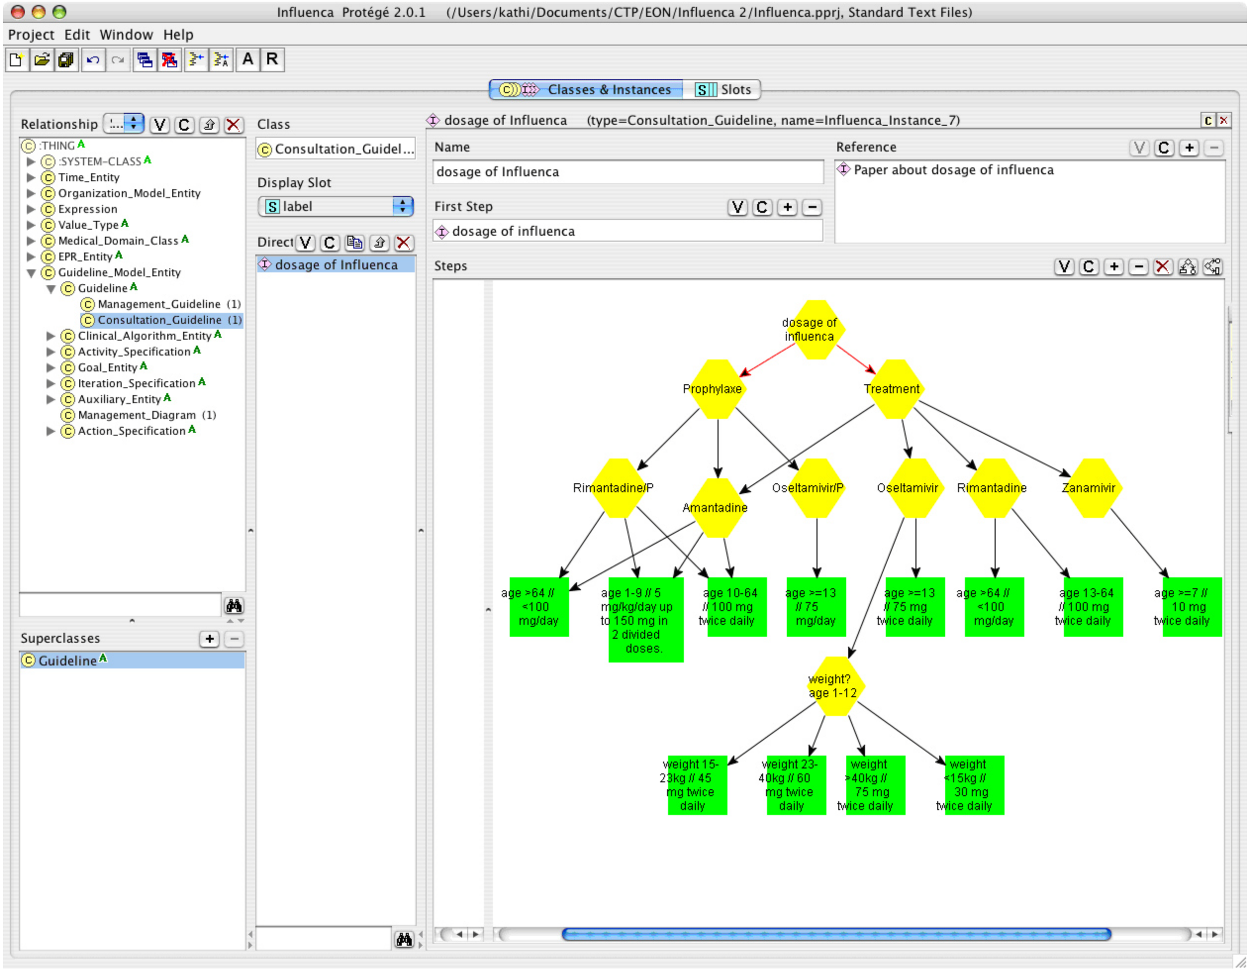

Influenca Protégé 2.0.1 (/Users/kathi/Documents/CTP/EON/Influenca 2/Influenca.pprj, Standard Text Files)

dosage of Influenca (type=Consultation_Guideline, name=Influenca_Instance_7)

Name: dosage of Influenca

Reference: Paper about dosage of influenca

First Step: dosage of influenca

Class: Consultation_Guidel...

Display Slot: label

Direct: dosage of Influenca

Superclasses: Guideline

Figure 2: Protégé [Gennari et al., 2002].

5.1.3. Tallis Toolset

Being the successor of Arezzo – the first software to generate, visualize, and execute PROforma [Fox et al., 1998] guidelines – the Tallis Toolset [Steele & Fox, 2002] is designed to create, publish and enact PROforma applications over the web. The PROforma task ontology is a set of reusable components: plans, decisions, actions, and enquiries (see Figure 3). These tasks are composed to form networks (called "workflow") and filled with details. All tasks have attributes like goals, preconditions, and postconditions.

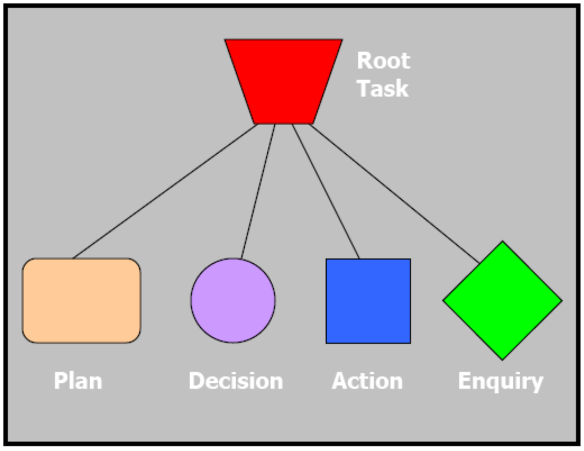

Figure 3: PROforma tasks [Steele & Fox, 2002].

The Tallis Toolset consists of three components: Composer, Tester, and Engine. The Composer provides a graphical user interface to visualize and support the generation of task networks (see Figure 4). The Tester is aimed at testing and debugging the logic of the generated networks, and the Engine executes the generated and tested PROforma guideline application.

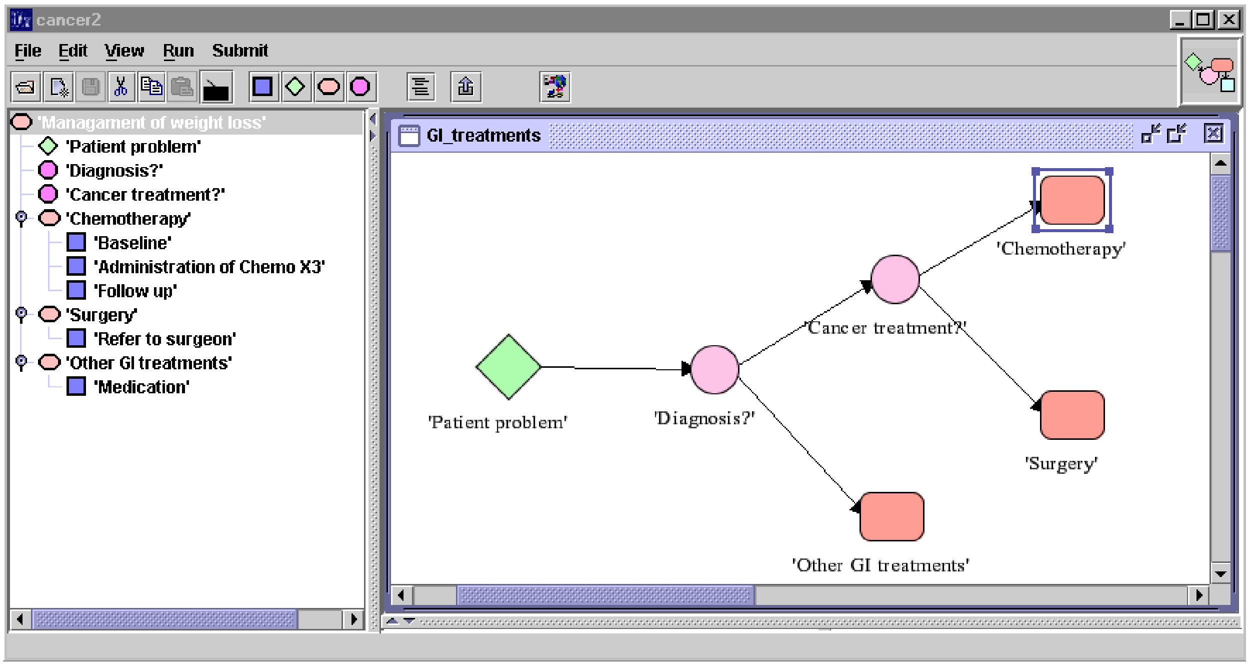

Figure 4: Tallis authoring environment: a simplified task network for cancer diagnosis and treatment [Steele & Fox, 2002].

5.1.4. GUIDE

GUIDE [Quaglini et al., 2001] was developed at the University of Pavia as part of a guideline modeling and execution framework. It serves a three-fold purpose, i.e., integrating CPGs into clinical workflow, using decision trees and influence diagrams to visualize complex relations, and using extended Petri Nets for the simulation of guideline implementation.
NEW-GUIDE [Ciccarese et al., 2003] is intended to facilitate the whole patient management, i.e., medical actions and organizational actions that can be managed by a workflow management system. The NEW-GUIDE editor allows for graphically authoring the workflow of guidelines in a flowchart-like representation (see Figure 5).

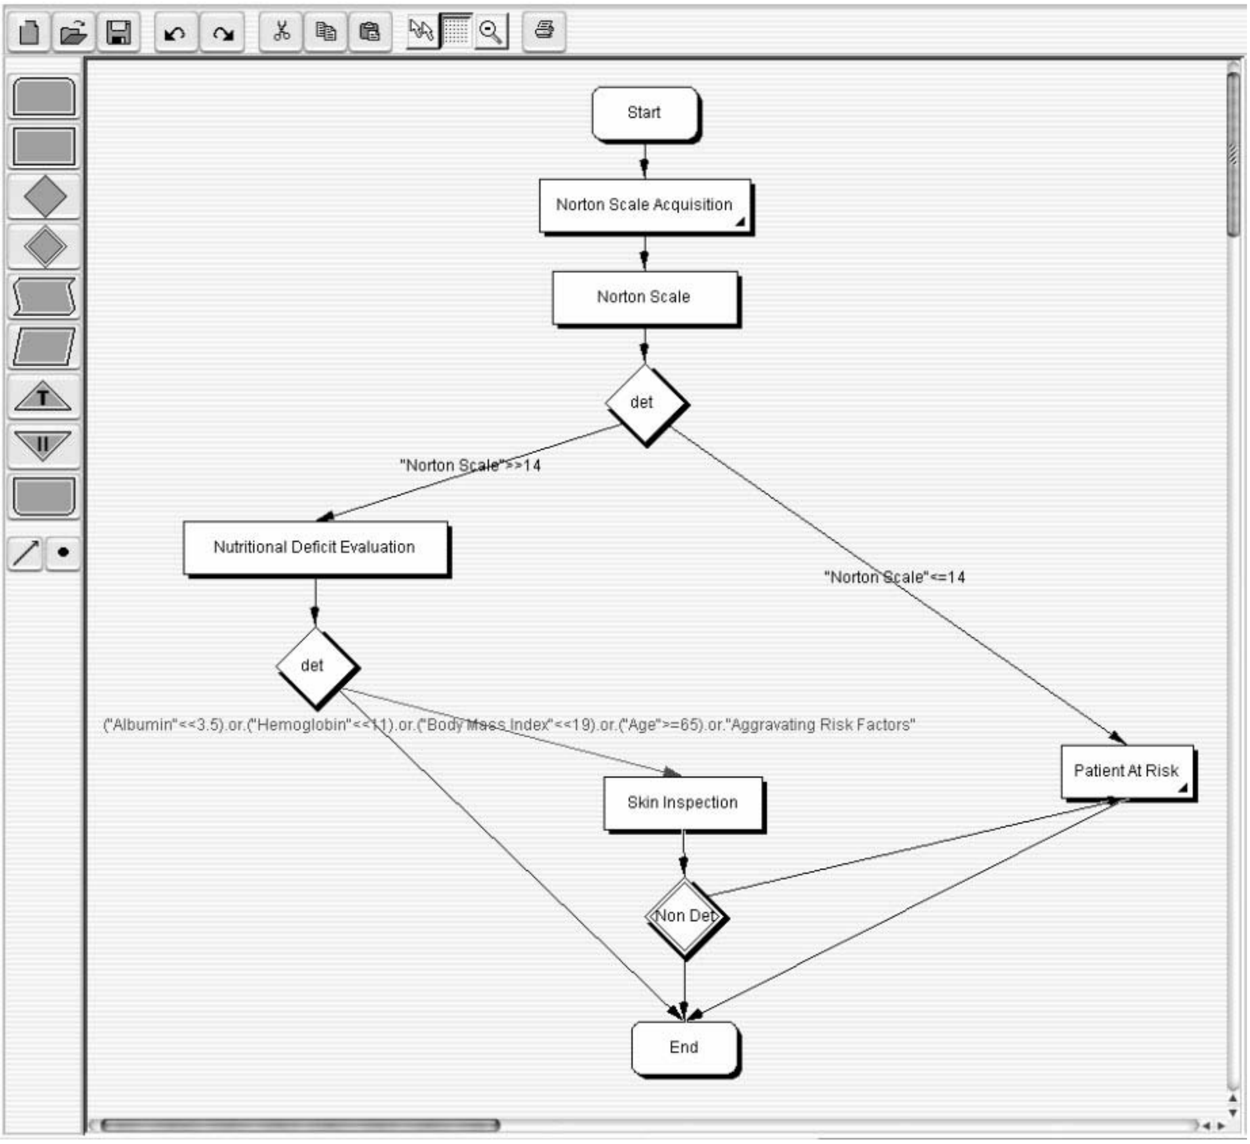
Start
Norton Scale Acquisition
Norton Scale
det
"Norton Scale">>14
Nutritional Deficit Evaluation
"Norton Scale"<=14
det
("Albumin"<<3.5).or.("Hemoglobin"<<11).or.("Body Mass Index"<<19).or.("Age">=65).or."Aggravating Risk Factors"
Patient At Risk
Skin Inspection
Non Det
End

Figure 5: The NewGuide editor [Ciccarese et al., 2004].

5.1.5. GLARE

GLARE [Terenziani et al., 2004] (GuideLine Acquisition, Representation, and Execution) is a domain independent tool. Similar to PROforma GLARE uses a limited set of representation primitives: work actions, query actions, decisions, and conclusions are called atomic actions. These primitives are composed to composite actions whereas relations between actions (e.g., sequence, alternative, or repetition) define the execution order. The Acquisition Graphical Interface (see Figure 6) provides primitives for drawing the guideline workflow similar to the representations of the Tallis Toolset and Protégé.

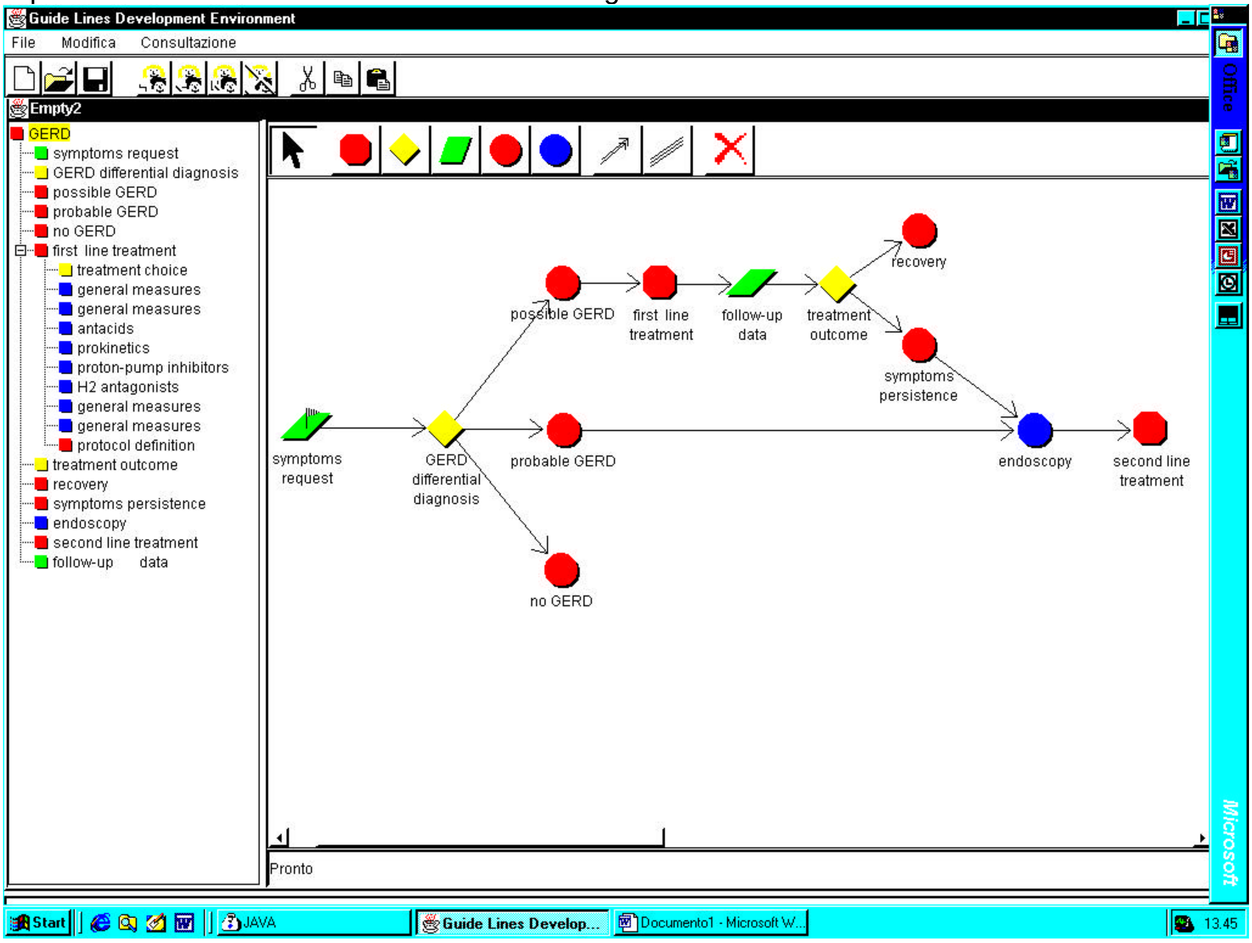

Figure 6: GLARE acquisition module graphical interface [Terenziani et al., 2004].

5.1.6. VisiGuide (Part of DeGeL)

VisiGuide is part of the DeGeL (Digital electronic Guideline Library) project [Shahar et al., 2003], which is a web-based architecture aimed at facilitating the transformation of a textual guideline into a formal model. The VisiGuide tool (see Figure 7) is used to browse guidelines and to visualize their structure. It supports the presentation of large amounts of guidelines organized by indexing semantic axes as well as the exploration of the different components of a single guideline.

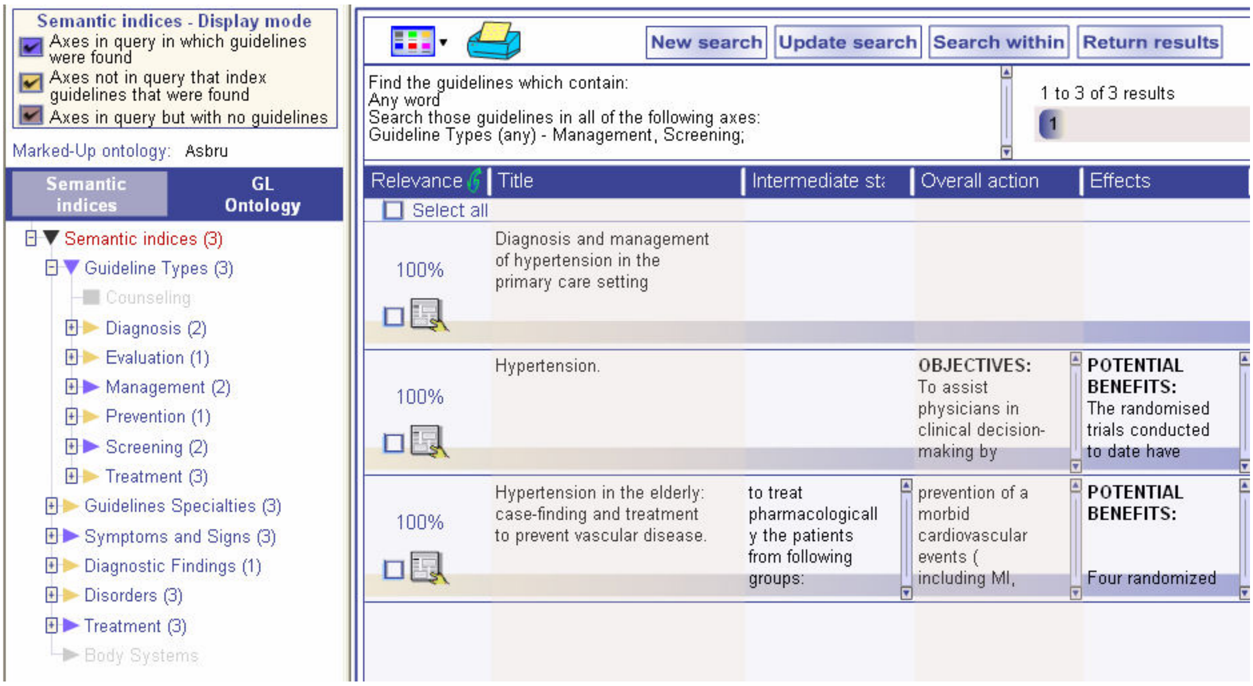

Find the guidelines which contain:
Any word
Search those guidelines in all of the following axes:
Guideline Types (any) - Management, Screening;

Relevance	Title	Intermediate st:	Overall action	Effects
100% □	Diagnosis and management of hypertension in the primary care setting			
100% □	Hypertension.		OBJECTIVES: To assist physicians in clinical decision-making by	POTENTIAL BENEFITS: The randomised trials conducted to date have
100% □	Hypertension in the elderly: case-finding and treatment to prevent vascular disease.	to treat pharmacologicall y the patients from following groups:	prevention of a morbid cardiovascular events (including MI,	POTENTIAL BENEFITS: Four randomized

Figure 7: VisiGuide [Shahar et al., 2003].

5.1.7. AsbruView

AsbruView [Kosara & Miksch, 2001] is part of the Asgaard/Asbru project [Shahar et al., 1998]. It is a graphical user interface to visualize the logical and temporal information of treatment plans expressed in the Asbru modeling language. AsbruView uses visual metaphors such as traffic signs and running tracks to communicate complex information, i.e., the hierarchical composition of plans, temporal order of plans, conditions, precise temporal constraints, temporal uncertainties, etc.

AsbruView consists of two parts: the Topological View and the Temporal View in order to serve different needs (see Figure 8).

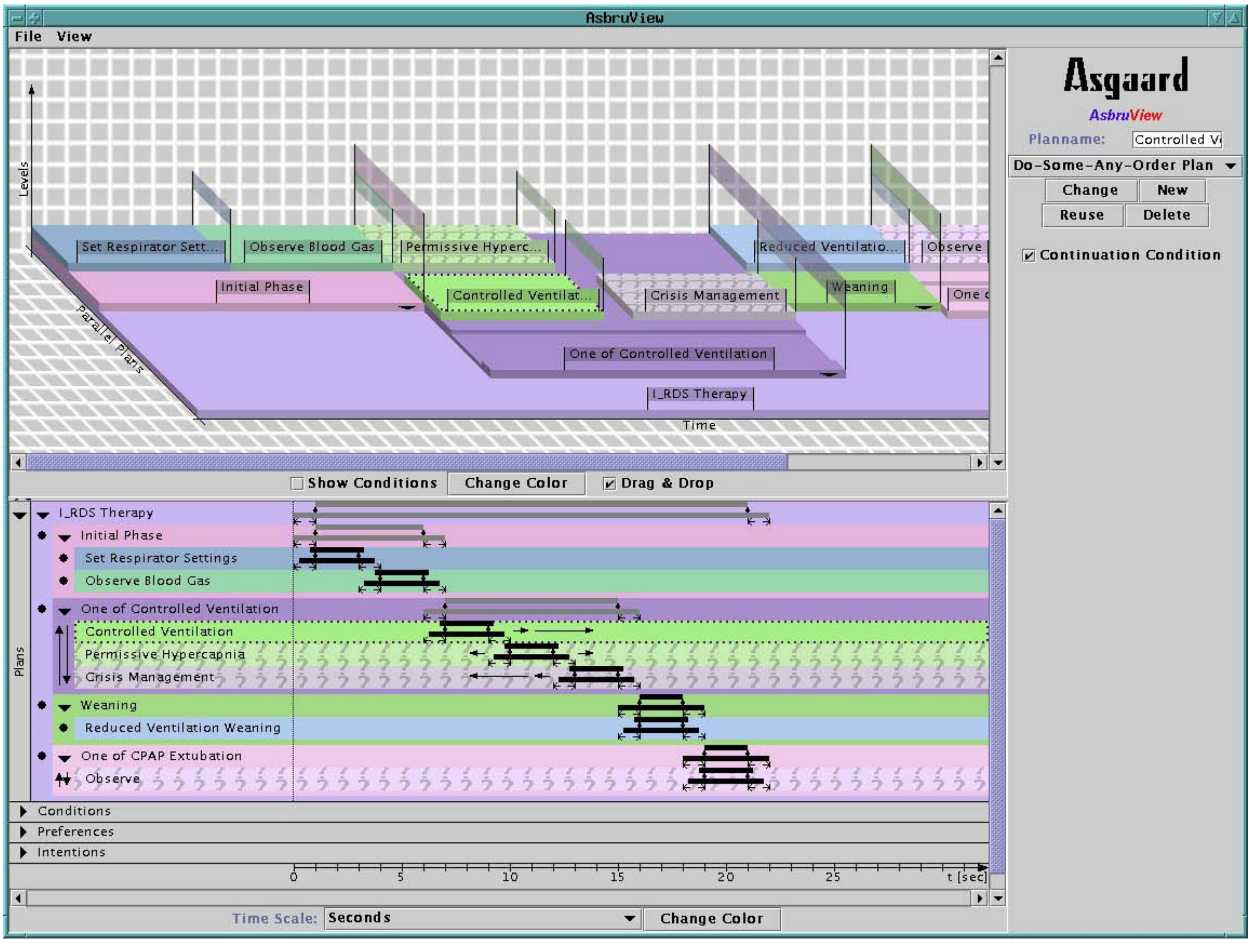

Planname: Controlled V

Do-Some-Any-Order Plan

☑ Continuation Condition

□ Show Conditions

☑ Drag & Drop

Time Scale: Seconds

Figure 8: AsbruView [Kosara & Miksch, 2001].

The **Topological View** is aimed at visualizing the sequence of plans and the relations between them, together with the different conditions of plans. The time axis does not give a precise time scale but serves to indicate the direction. The central metaphor of the Topological View is a running rack representing an Asbru plan – the physician is considered to run along this track while treating the patient. Stacking plans on top of each other indicates the hierarchical decomposition. The sequential order (e.g., plan A happens after plan B, plan A starts at the same time as plan B) is communicated by laying running track plan objects out along or parallel to the time axis. Sub-plans of a sequential plan are put on the containing plan after each other parallel to the time axis. Sub-plans of a parallel plan are put on the containing plan next to each other along the parallel plans axis. Sub-plans of unordered plans are put on the containing plan in a pattern that is meant to show that it has nothing to do with the real execution sequence. Cyclical plans are indicated by a loop-arrow from the end of the running track back to the beginning. The visualization of any-order plans is not supported by AsbruView. Optional plans are marked by a question-mark texture.

Again, graphical metaphors are used to visualize conditions of plans, i.e., metaphors from traffic control (see Figure 9). To visualize the **filter precondition** of a plan AsbruView uses the “no entrance with exceptions” sign. A barrier is used for the **setup precondition,** which is opened as soon as the patient meets the criteria. A traffic light is used to indicate three conditions: the red light stands for the **abort condition**, the yellow light is used for the **suspend condition**, and the green light stands for the **reactivate condition** when after suspension the criteria for reactivating the plan are met again. The plan’s **complete condition** is visualized by a finishing flag. When a condition is not defined for a plan, the corresponding traffic sign is drawn in gray.

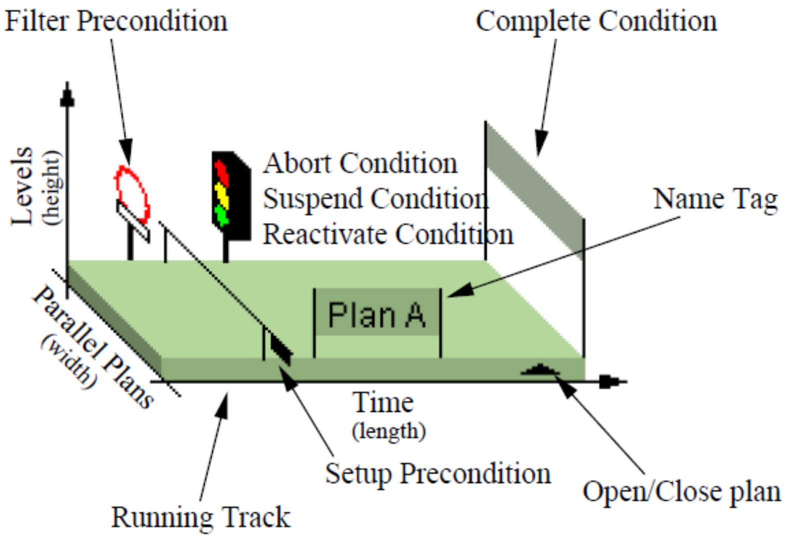

Figure 9: AsbruView: Anatomy of an Asbru Plan in the Topological View [Kosara & Miksch, 2001].

The **Temporal View** is based on time lines. Sub-plans are drawn indented and inside the rectangle representing the containing plan which indicates the hierarchical decomposition. The type of plans, i.e. the execution order of their sub-plans, is indicated by small symbols left of the sub-plans. However, the Temporal View is focused on the communication of the exact temporal dimensions of the plans. This is accomplished by using time-glyphs. A time-glyph consists of two bars; the first bar indicates the maximum duration (MaxDur) of the plan. Beneath this bar the earliest starting shift (ESS), latest starting shift (LSS), earliest finishing shift (EFS), and latest finishing shift (LFS) are indicated by little arrows pointing towards these time points. On top of the MaxDur bar two diamonds represent the LSS and the EFS, on which the second bar rests representing the minimum duration (MinDur) of the plan. If the LSS or the EFS are not defined, it is represented by a roll instead of a diamond to indicate that it might move.

This time-glyph implies that the MinDur bar and the MaxDur bar cannot be shorter than the distance between the LSS and the EFS; otherwise the MinDur bar would fall down between the diamonds or the diamonds would fall down the MaxDur bar. Parts, which are not defined, are drawn in gray again.

The time-glyph also communicates information about the granularity of time annotations. If the same granularity is used for all parts, they all have sharp ends. If one time point has a finer granularity than the current display, the time point is represented by a circle in analogy to open intervals in mathematics (e.g. “from 0 to, but not including 1”). If the time point has a coarser granularity than the current display, it is represented by a zigzag pattern (see Figure 10).

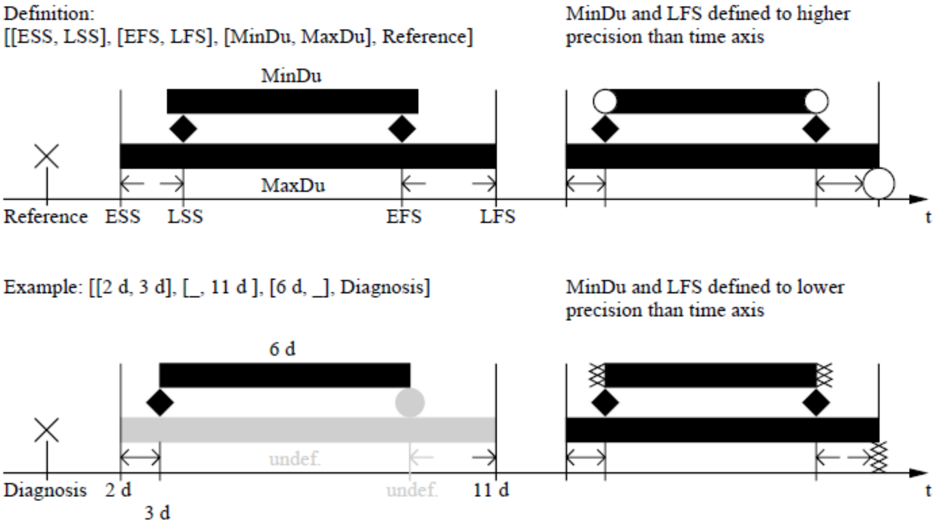
Definition:
[[ESS, LSS], [EFS, LFS], [MinDu, MaxDu], Reference]
MinDu
MaxDu
Reference
ESS
LSS
EFS
LFS
t
MinDu and LFS defined to higher precision than time axis
Example: [[2 d, 3 d], [_, 11 d], [6 d, _], Diagnosis]
6 d
undef.
Diagnosis
2 d
3 d
undef.
11 d
t
MinDu and LFS defined to lower precision than time axis

Figure 10: Time Annotations [Kosara & Miksch, 2001].

5.2. Visualization of a Patient’s Condition

In the following subsections we outline selected approaches to visualize patient data.

5.2.1. Graphical Summary of Patient Status

The Graphical Summary of Patient Status [Powsner & Tufte, 1994] uses small repeated graphs with identical formats to map different finidings and treatments over time. In doing so, they managed to graphically summarize the data of a 5cm thick record on a single page. A non-linear time-scale compresses years of data into a context for assessing recent trends. Scaled values are plottet on a vertical axis whereas five same sized qualitative ranges of normal, elevated, reduced, critically elevated, and critically reduced values are indicated (see Figure 11 and Figure 12). The Graphical Summary of Patient Status makes use of concepts like small multiples, Focus+Context, or the integration of textual and graphical information.

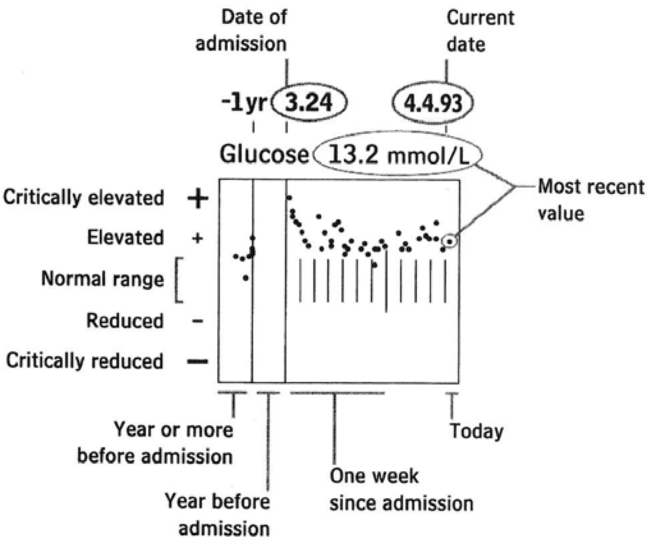

Figure 11: Graphical Summary of Patient Status – graph [Powsner & Tufte, 1994].

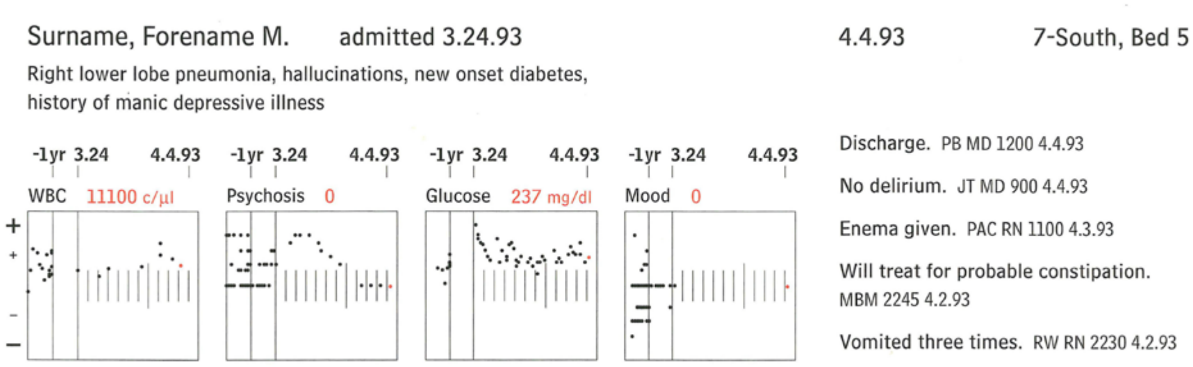
Surname, Forename M. admitted 3.24.93 4.4.93 7-South, Bed 5

Right lower lobe pneumonia, hallucinations, new onset diabetes, history of manic depressive illness

Discharge. PB MD 1200 4.4.93

No delirium. JT MD 900 4.4.93

Enema given. PAC RN 1100 4.3.93

Will treat for probable constipation. MBM 2245 4.2.93

Vomited three times. RW RN 2230 4.2.93

Figure 12: Part of the summary produced by the Graphical Summary of Patient Status [Powsner & Tufte, 1994].

5.2.2. Time Lines and LifeLines

A very intuitive way of representing the temporal aspects of an activity is drawing a horizontal line on a time scale. This method is called Time Lines which are horizontal lines or bars depicting time spans and sequences; they are a powerful visualization method for time-oriented data. LifeLines [Plaisant et al., 1998] is a software tool for visualizing the time oriented data of personal history records on a zoom-able timeline (see Figure 13). They use Time Lines with some extensions, for instance so-called “facets” are introduced for grouping aspects of the record (e.g. problems, allergies, diagnosis, labs, imaging, medications, immunizations, etc.) which can be expanded and collapsed. It was applied to summarize medical records as a set of lines and events. Moreover, LifeLines uses the color (e.g. a parameter with a status marked as "alert" or severe allergies are red) and height (e.g. significance of action or event) of individual bars to encode information.

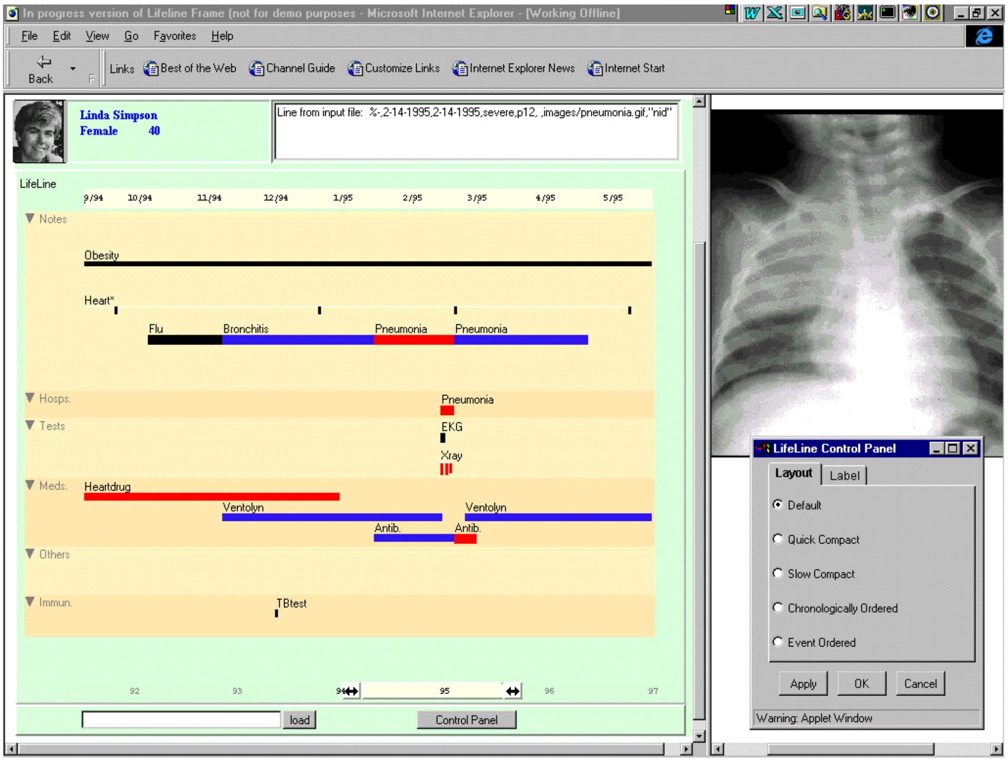

In progress version of Lifeline Frame [not for demo purposes - Microsoft Internet Explorer - [Working Offline]

Linda Simpson
Female 40

Line from input file: %-,2-14-1995,2-14-1995,severe,p12, ,images/pneumonia.gif,"nid"

LifeLine Control Panel

Layout: ● Default

Warning: Applet Window

Figure 13: LifeLines used to visualize a patient record; zoomed to reveal details of a long pneumonia episode [Plaisant et al., 1998].

LifeLines2 [Wang et al., 2008] is an extension of LifeLines, which was designed to visualize subsets of medical records from multiple patients, for instance visualizing all records of patients with pneumonia and influenza. LifeLines2 provides special features to align, rank, and filter the results of queries. These features facilitate the efficient comparison of different patient records and thus, help to discover previous unseen patterns (see Figure 14).

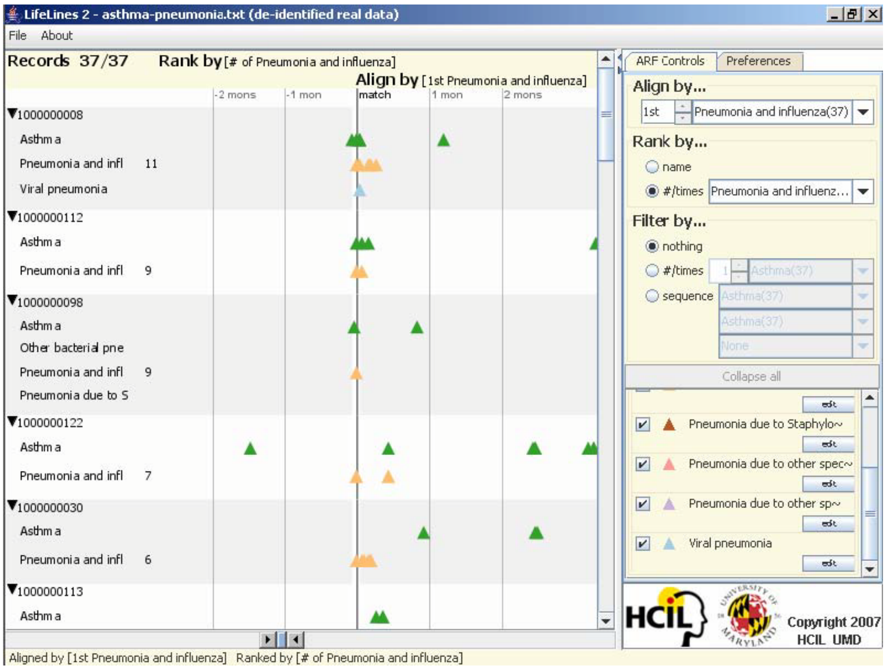

LifeLines 2 - asthma-pneumonia.txt (de-identified real data)

Records 37/37 Rank by [# of Pneumonia and influenza]

Align by [1st Pneumonia and influenza]

Align by... 1st Pneumonia and influenza(37)

Rank by... ○ name ● #/times Pneumonia and influenz…

Filter by... ● nothing ○ #/times 1 Asthma(37) ○ sequence Asthma(37) Asthma(37) None

☑ Pneumonia due to Staphylo~

☑ Pneumonia due to other spec~

☑ Pneumonia due to other sp~

☑ Viral pneumonia

Aligned by [1st Pneumonia and influenza] Ranked by [# of Pneumonia and influenza]

Copyright 2007 HCIL UMD

Figure 14: LifeLines2: All patient records are aligned by the 1st *Pneumonia and influenza* along a relative time scale. This alignment allows for an efficient identification of the diagnosis of asthma for the first 3 patients within a month prior or at the same time their pneumonia was diagnosed [Wang et al., 2008].

5.2.3. PatternFinder

The PatternFinder query tool [Plaisant et al., 2008] has been developed to facilitate the specification of queries in SQL and subsequently visualize the results of the query. It is implemented in a large operational system at the Washington Hospital Center. By providing means to easily specify SQL queries - for instance, “Find patients who had a normal serum creatinine lab test less than 2 days before a radiology test with intravenous contrast, followed by an increase in serum creatinine by more than 50% and of more than 1.0 mg/dl within 5 days after the contrast administration.” – and by providing a visual representation of the results, the PatternFinder allows for the analysis of temporal patterns in patient histories. The user interface is divided into three panes (see Figure 15): the filter pane at the top left contains widgets to specify the query (i.e., the temporal scope, the sentinel event, the baseline event, and the follow on event, etc.). The results of the query are provided by a table in the bottom left pane; in addition, the pane at the right displays a visual representation of the results. A LifeLine [Plaisant et al., 1998] is drawn for each patient found and all events that match the query are indicated. The user can expand the LifeLines for more detailed information about the events.

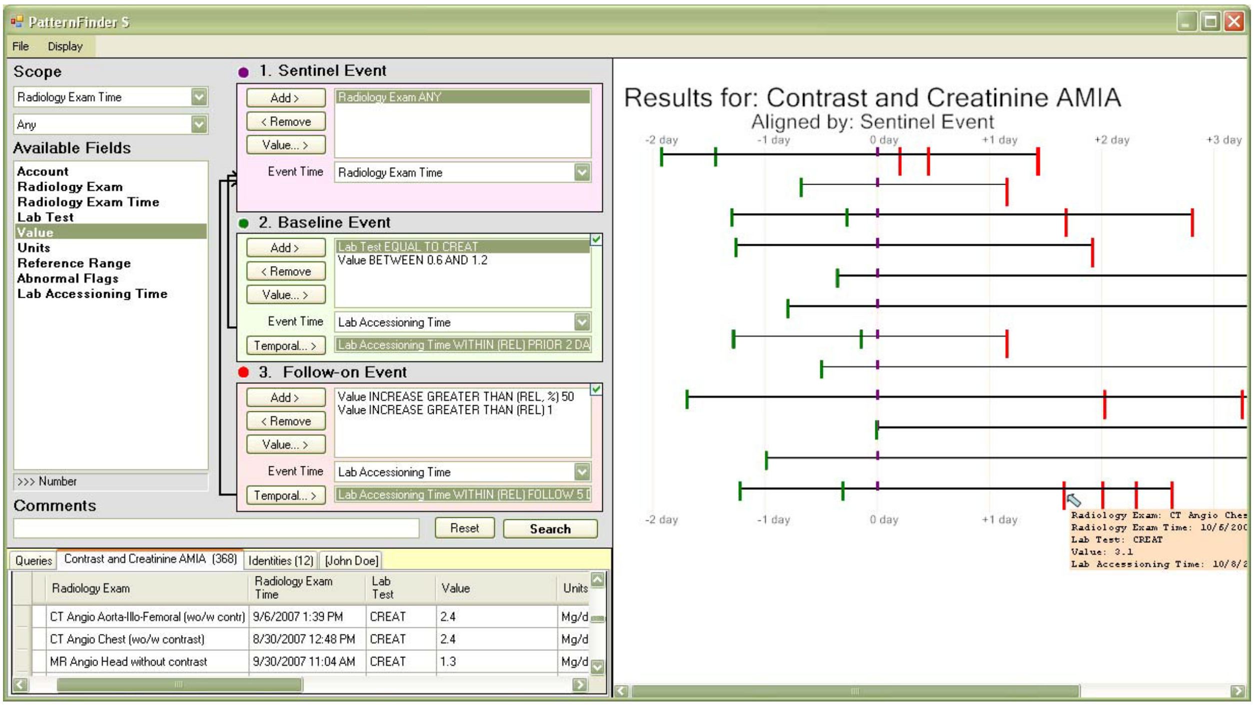

Scope: Radiology Exam Time; Any

1. Sentinel Event: Radiology Exam ANY
Event Time: Radiology Exam Time

2. Baseline Event: Lab Test EQUAL TO CREAT
Value BETWEEN 0.6 AND 1.2
Event Time: Lab Accessioning Time
Temporal: Lab Accessioning Time WITHIN (REL) PRIOR 2 DA

3. Follow-on Event: Value INCREASE GREATER THAN (REL, %) 50
Value INCREASE GREATER THAN (REL) 1
Event Time: Lab Accessioning Time
Temporal: Lab Accessioning Time WITHIN (REL) FOLLOW 5 I

Radiology Exam	Radiology Exam Time	Lab Test	Value	Units
CT Angio Aorta-Illo-Femoral (wo/w contr)	9/6/2007 1:39 PM	CREAT	2.4	Mg/d
CT Angio Chest (wo/w contrast)	8/30/2007 12:48 PM	CREAT	2.4	Mg/d
MR Angio Head without contrast	9/30/2007 11:04 AM	CREAT	1.3	Mg/d

Radiology Exam: CT Angio Ches
Radiology Exam Time: 10/6/200
Lab Test: CREAT
Value: 3.1
Lab Accessioning Time: 10/8/2

Figure 15: PatternFinder provides means to define a query in SQL, for instance: "Find patients who had a normal serum creatinine lab test less than 2 days before a radiology test with intravenous contrast, followed by an increase in serum creatinine by more than 50% and of more than 1.0 mg/dl within 5 days after the contrast administration". For each patient found, a LifeLine with marks for matching events is drawn [Plaisant et al., 2008].

5.2.4. KNAVE and KNAVE-II

KNAVE [Shahar & Cheng, 1998] has been developed for the interpretation, summarization, visualization, explanation, and exploration of time-oriented clinical data. KNAVE-II [Shahar et al., 2006] is en extension of KNAVE which supports the interpretation of time-oriented data as well as the visual exploration of the data at multiple levels of abstraction. The data can be represented along an absolute time scale as well as along a relative time scale. The interface contains a browser to the clinical domain's ontology on the left hand side. To investigate the data of one patient, the user may select raw data or an abstract concept from the ontology tree or s/he may ask a query from the left hand bottom panel. For each query, the found concepts are visualized in a separate panel at the right hand side of the interface. Quantitative data is represented by line charts while qualitative data is represented by qualitative regions (see Figure 16). Moreover, interactions for synchronizing the displays of the panels according to time, adding statistics to the panels, zooming in and out, etc., are provided.

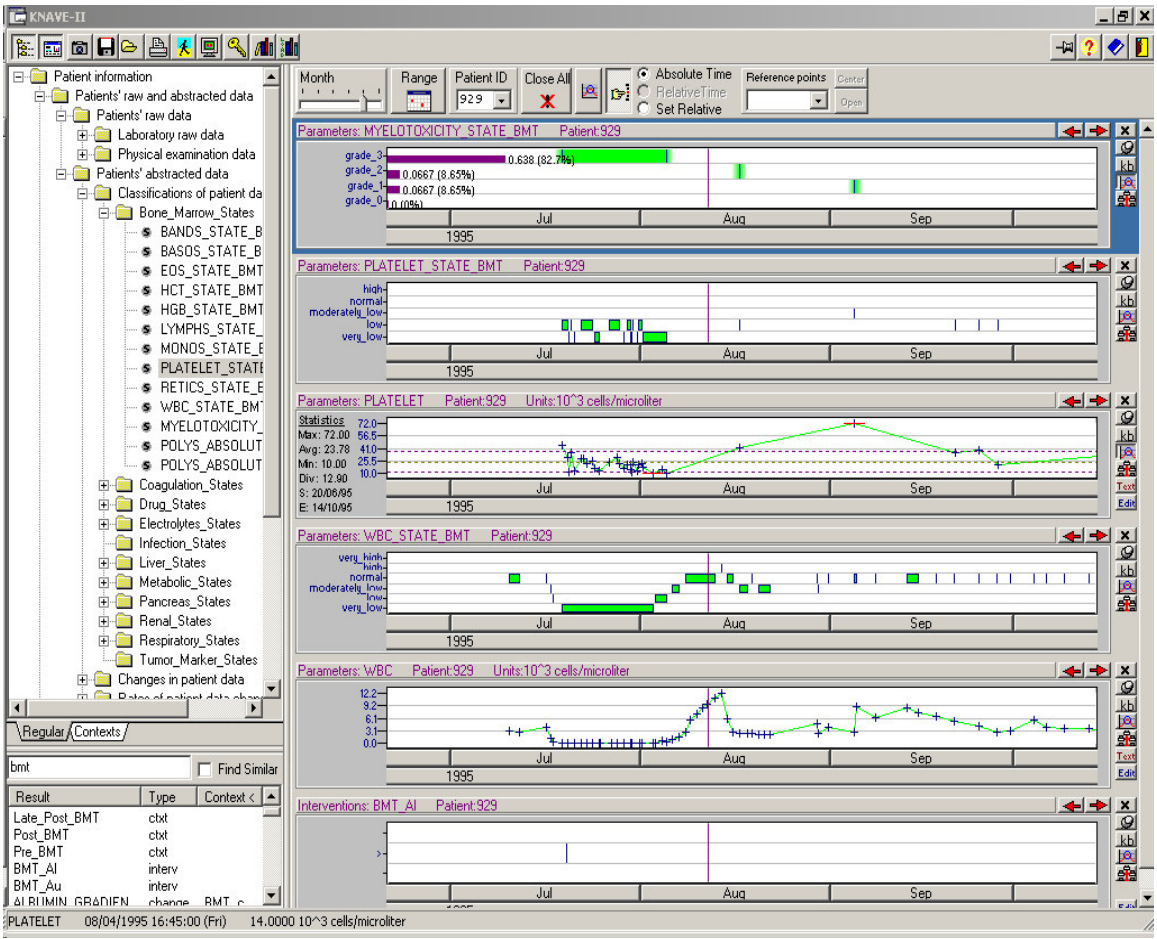

Statistics
Max: 72.00
Avg: 23.78
Min: 10.00
Div: 12.90
S: 20/06/95
E: 14/10/95

Patient ID: 929

bmt

Result	Type	Context (
Late_Post_BMT	ctxt	
Post_BMT	ctxt	
Pre_BMT	ctxt	
BMT_AI	interv	
BMT_Au	interv	
ALBUMIN_GRADIEN	change	BMT_c

PLATELET 08/04/1995 16:45:00 (Fri) 14.0000 10^3 cells/microliter

Figure 16: KNAVE-II [Shahar et al., 2006]: different parameters of one patient can be explored supported by multiple features, like different abstraction levels, absolute and relative time scales, different granularities of time, additional statistics, etc.

5.2.5. VISITORS

VISITORS [Klimov & Shahar, 2005] has been designed to allow for the exploration of time-oriented parameters of multiple patient records. Special attention has been paid to the visualization of temporal aspects, such as different time-granularities. In addition, raw data values can be represented as well as abstract data. Figure 17 shows the visualization of one parameter of raw data for a group of patients, while Figure 18 shows one parameter of qualitative values for a group of patients. Figure 19 shows the visual comparison of multiple parameters of raw data and abstract data. Interactions like Zoom+Pan, filtering of data, and details on demand facilitate the exploration process.

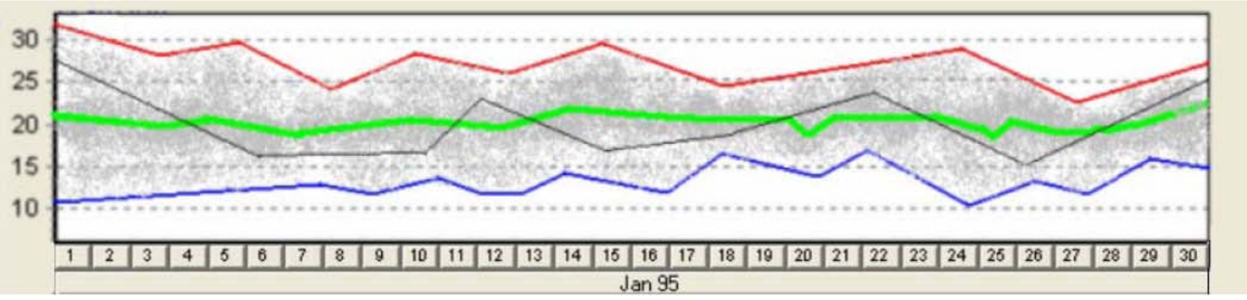

Figure 17: VISITORS [Klimov & Shahar, 2005]: visualization of a single raw parameter for a group of patients over time. The minimum and the maximum value of the group are indicated by the bottom line and the top line; the average value is represented by the green wide line in the middle. The gray line represents the data of one selected patient.

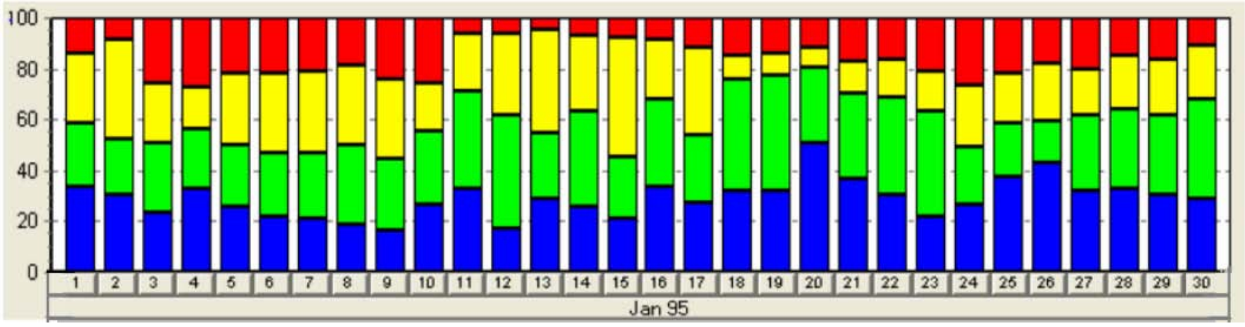

Figure 18: VISITORS [Klimov & Shahar, 2005]: visualization of one abstract parameter for a group of patients over time. The qualitative value is encoded by colour. The height of the bars indicates the percentage of patients of the selected group with respect to the qualitative parameter.

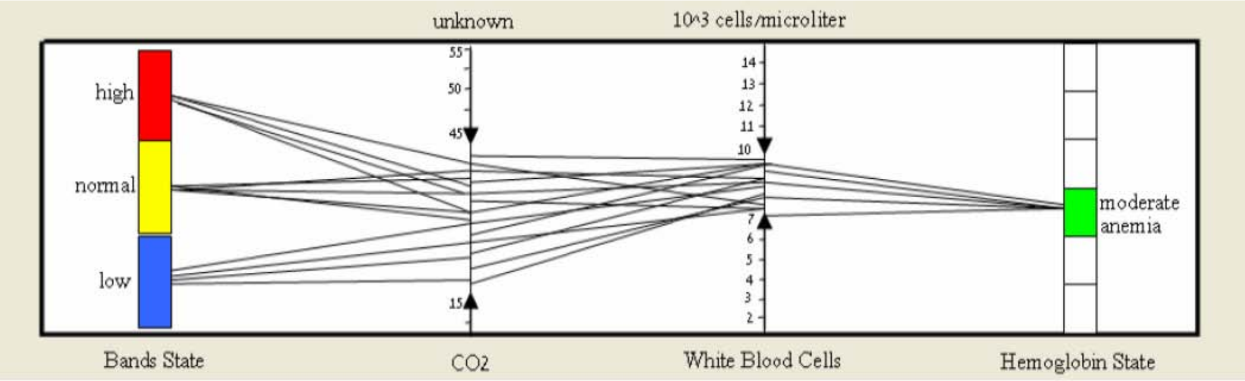

Figure 19: VISITORS [Klimov & Shahar, 2005]: visualization of multiple parameters (raw and abstract) for a group of patients. The units of the corresponding axes are displayed at the top of each axis. Again, qualitative values are indicated by colour and additionally, in textual form. For each parameter the mean values of one patient (during a selected time period) are connected by lines.

5.2.6. VIE-VISU

VIE-VISU [Horn et al., 1998] uses small interactive glyphs (small multiples) to represent the change in the patient's condition (in particular neonatal intensive care data) over time. The course of the patient is depicted by 24 multiples, whereas different geometric and visual attributes of one glyph encode important parameters from circulation, ventilation, and fluid balance, thus combining 15 different patient's parameters (see Figure 20). The motivation for this representation was to find a visualization supporting the immediate identification of stable patients, continuous improvement or worsening, as well as drastic changes in the patient's status.

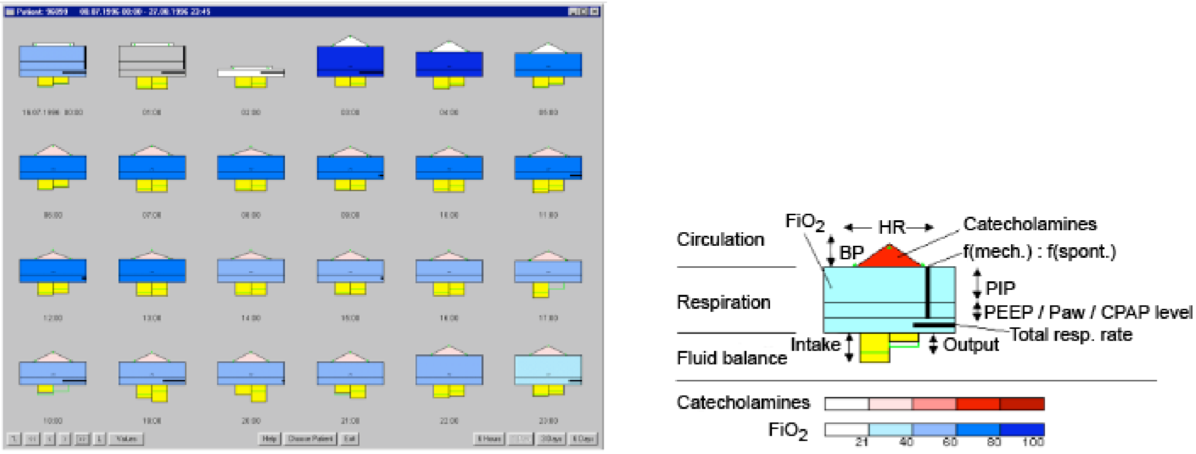

Figure 20: 24 hours display of VIE-VISU along with the legend of the glyph [Horn et al., 1998].

5.2.7. Interactive Parallel Bar Charts (IPBC)

Interactive Parallel Bar Charts have been designed by [Chittaro et al., 2003] in order to replace tens of separate screens used in traditional hemodialysis systems. Along with each hemodialysis session a time series of values is produced for each recorded clinical parameter. IPBC represent such a time series in a bar chart format, whereas the X axis represents the time and the Y axis represents the recorded value (the height of the bar). A third dimension is used to layout the time series of different hemodialysis sessions next to each other (see Figure 21).
In terms of supporting the interaction with the representation, several useful features are provided. In order to overcome typical problems with 3D visualizations, like a long mouse movement from the menu bar to the visualization or the waste of screen space towards the corners of the screen, a RoundToolbar was introduced which can be positioned next to the visualization. In addition, they provide pre-defined movements along limited trajectories to ease the navigation in 3D space. The user can define the color-coding to classify time series values, interactive features to move and modify the representation are provided, and pattern matching can be performed with configurable tolerance levels. Addressing a common problem with 3D representations, i.e. the occlusion of elements, Chittaro et al. provide a possibility to flatten occluding elements.

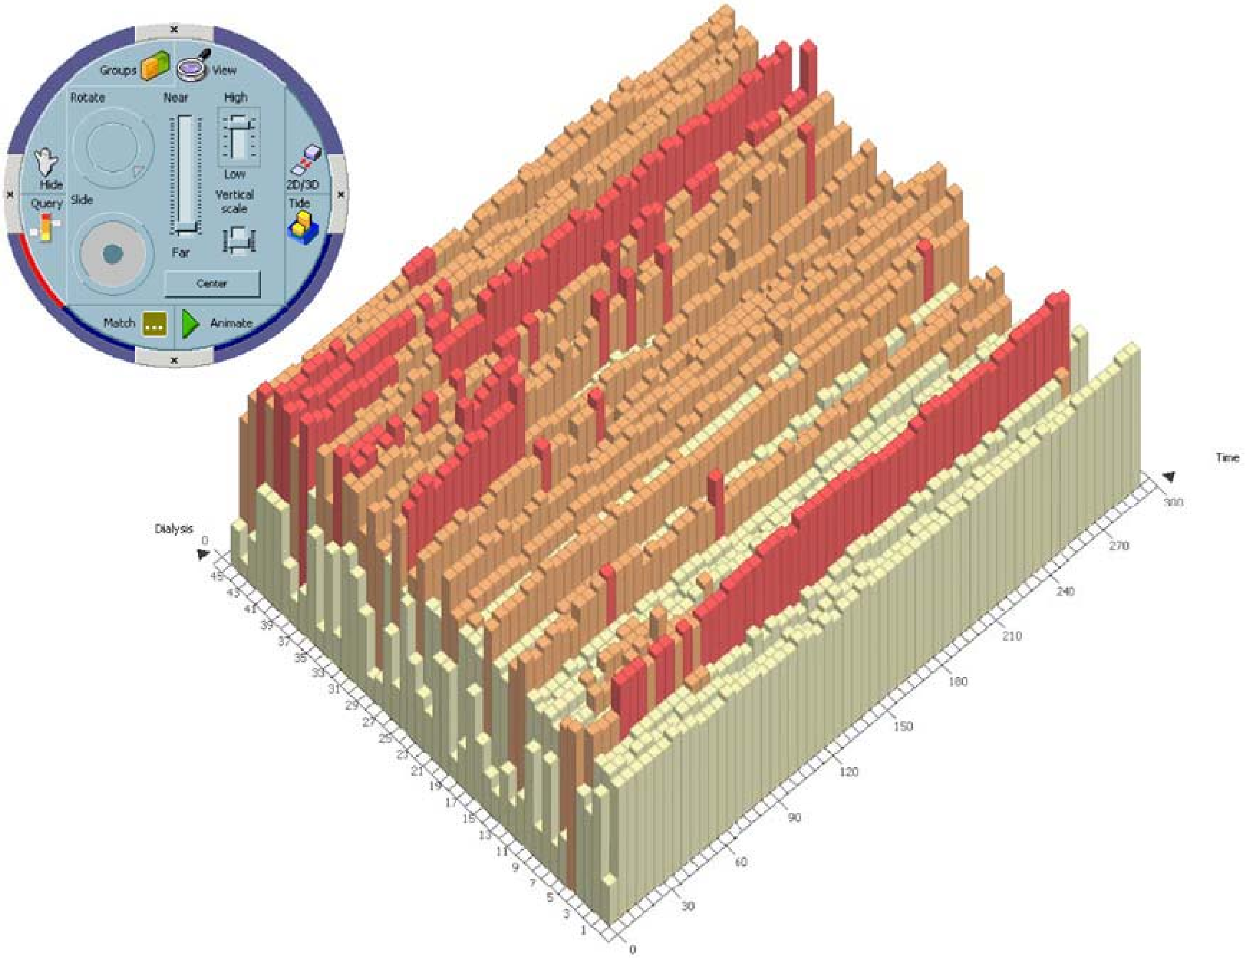

Figure 21: Interactive Parallel Bar Charts [Chittaro et al., 2003]: the X axis (at the right) represents the time, the Y axis (the height of the bar) represents the recorded values of one parameter during a hemodialysis session, and the Z axis (at the left) shows the different hemodialysis sessions.

5.2.8. Gravi++

Gravi++ [Hinum et al., 2005] has been designed to visualize an extensive set of questionnaires, which need to be analyzed to capture the therapeutic process of anorectic girls and to find predictors. Psychologists and physicians need to derive time-oriented quantitative and qualitative information and the states of the patients. Gravi++ uses two kinds of icons to represent patients and questions for the questionnaires. The questions' icons are arranged in a circle. A spring based system is used to arrange patient's icons within this circle as follows: each person is connection with each question by an invisible spring. The pulling strength of the individual spring depends on the answer of the patient to this question, i.e. the higher the question was answered (on a scale from "0" to "6") the stronger the person's icon is pulled near. Thus, the answers the patient gave in the questionnaire influence the position of the patient's icon on the screen, which leads to the formation of clusters of persons who gave similar answers (see Figure 22).

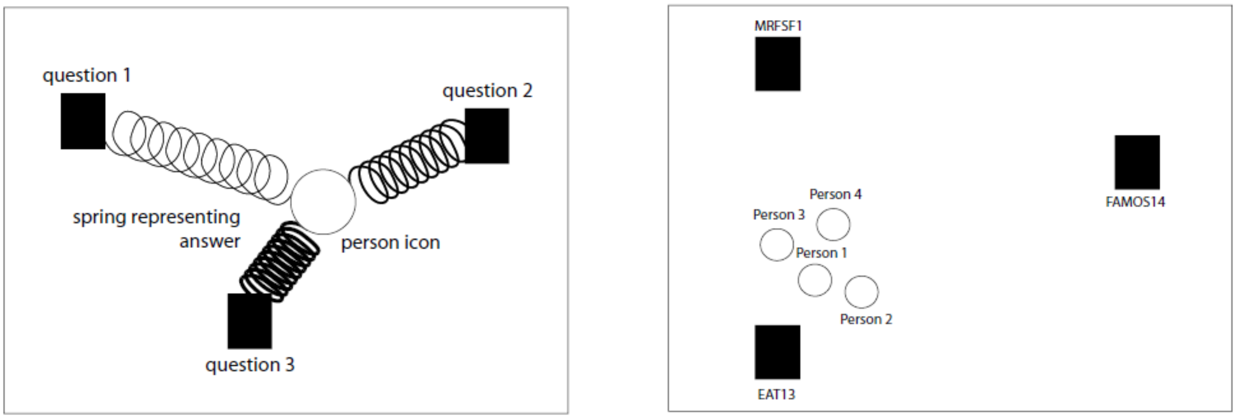

Figure 22: Spring based visualization: A person is connected with each question by (invisible) springs. The higher the person answered a question, the stronger the corresponding spring pulls the person's icon near. This leads to the formation of clusters of persons who gave similar answers [Hinum et al., 2005].

Moreover, the size of a person's icon can be mapped to an additional parameter (e.g., to the body mass index of the patient) or to the attraction force. To visualize the change of the patients over time, Gravi++ uses animation or traces. Rings around the questions' icons along with Star Glyphs are used to communicate exact values for the attraction of a patient to the corresponding question. The Star Glyphs are connected with the question rings; both are drawn in the color of the corresponding person in order to facilitate the identification of visual elements that belong together (see Figure 23).
Interactive features, for instance changing the position or strength of a question or highlighting a question or a person, facilitate the exploration of the data.

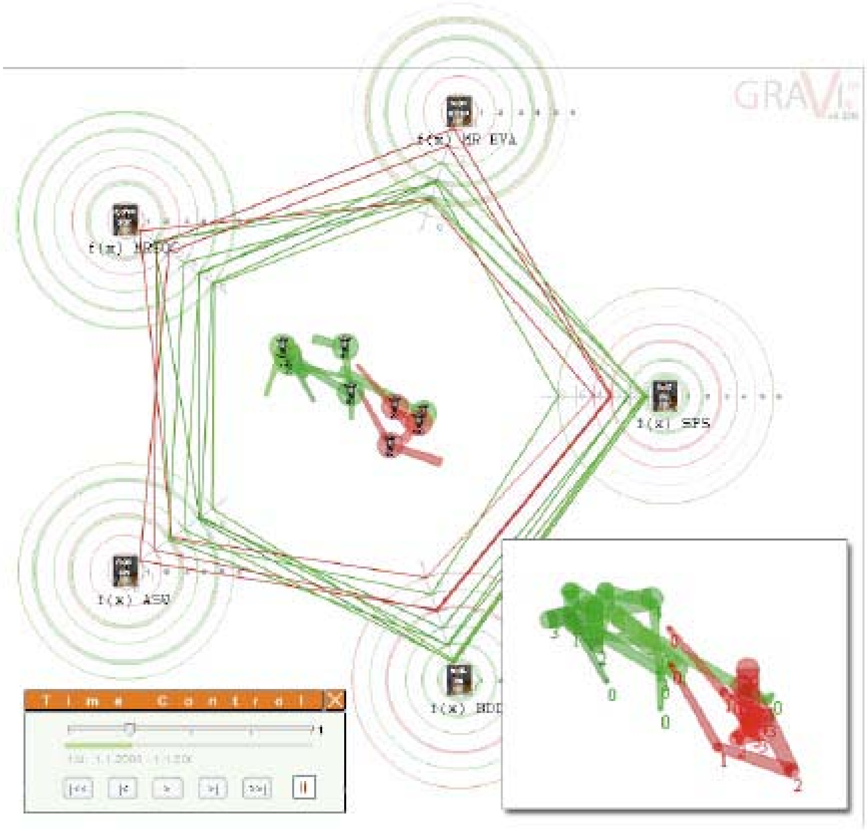

Figure 23: Gravi++ [Hinum et al., 2005]: the course of these two patients over time is visualized by animation or by traces (as shown in detail on the lower right).

5.3. Visualization of Treatment Plans in Connection with a Patient's Condition

A further visualization challenge is to present a patient's data along with his/her treatment process. Combining the visualization of treatment plans and the visualization of a patient's condition is an important means for physicians, but only few approaches exist that deal with this requirement.

5.3.1. Guideline Overview Tool (GOT)

The main purpose of the Guideline Overview Tool [Aigner, 2001] is to provide a compact overview of patient data in connection with applied treatment plans as well as to allow for the comparison of the course of several patients at one view. GOT uses an extension of LifeLines [Plaisant et al., 1998] to visualize the various states of Asbru plans as well as their execution order. Moreover, plans in the future are visualized as non-filled bars with dotted edges.
Qualitative patients' parameters (e.g., no, moderate, and high fever) are visualized by color coded LifeLines while quantitative data is represented by diagrams over time. Due to restricted vertical space, at most three parameters are displayed at once. The user interface of GOT is divided into three panes: the time scale can be found at the top of the interface; the biggest pane shows treatment plans and patients' parameters and is located beneth the time scale. Additional data about the patients (variables, constants, and current values) is displayed at the right pane (see Figure 24).

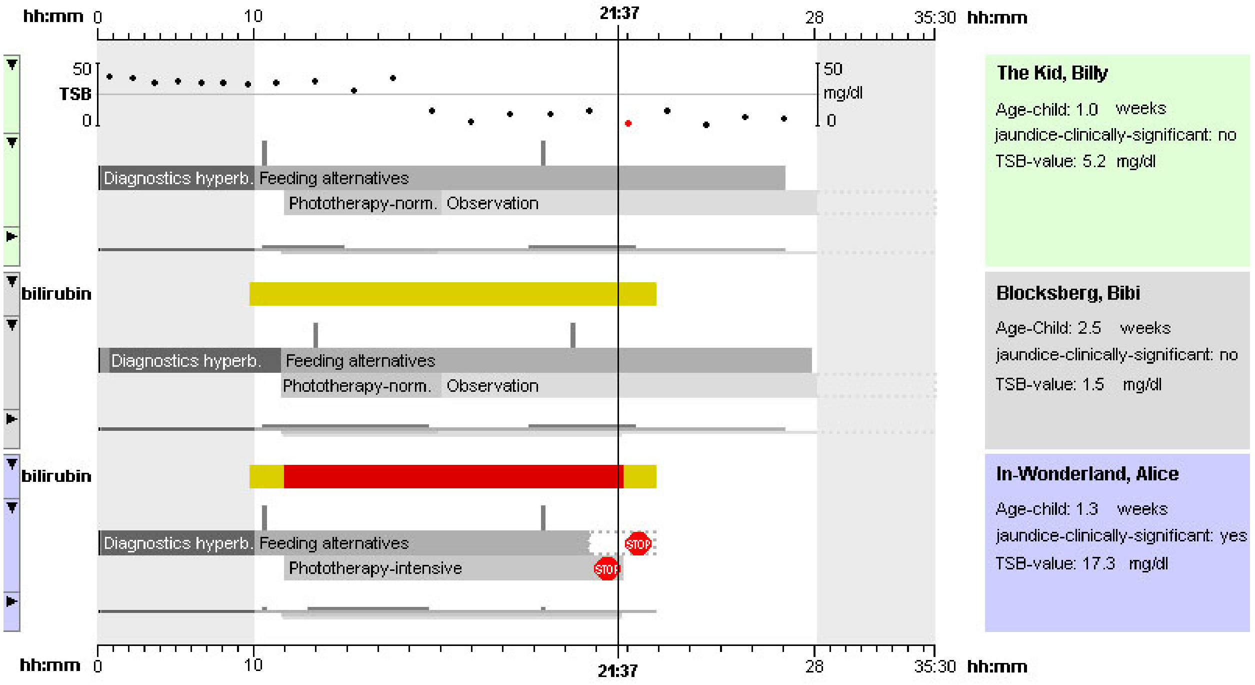

The Kid, Billy
Age-child: 1.0 weeks
jaundice-clinically-significant: no
TSB-value: 5.2 mg/dl

Blocksberg, Bibi
Age-Child: 2.5 weeks
jaundice-clinically-significant: no
TSB-value: 1.5 mg/dl

In-Wonderland, Alice
Age-child: 1.3 weeks
jaundice-clinically-significant: yes
TSB-value: 17.3 mg/dl

Figure 24: Guideline Overview Tool: visualization of the course of three patients treated with the same guideline for Hyperbilirubinemia [Aigner, 2001].

5.3.2. Midgaard

Midgaard [Bade et al., 2004] is an interactive visualization application for intensive care data. It combines quantitative representations of data with qualitative representations, which is called hybrid technique. A quantitative data stream, for instance a fever curve, is enhanced with colour-coded qualitative regions which add additional knowledge to the presentation, for instance if a particular value is critical or not (see Figure 25).

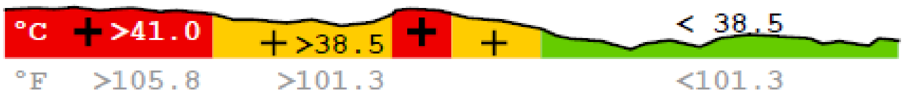

Figure 25: Qualitative - quantitative hybrid as used in Midgaard. The quantitative data of a fever curve is represented by the height while additional qualitative information is provided by the color-coding (critical fever values are colored red, moderate fever yellow and normal temperature green) [Bade et al., 2004].

In [Bade et al., 2004] an other visualization technique is introduced to overcome the problem of very limited space. When displaying high-frequency data many data values would have to be depicted at the same small display region (e.g. a pixel or a row of pixels). Thus, they use a redesign of *Tukey box-plots* – inspired by [Tufte, 1983]. The space between the minimum and the maximum represented by a light color and the space between the 25% and 75% percentile is colored a little bit darker; the median is marked black. Now, a lot of these small box-plot-bars can be connected to a full data stream representation on very limited display space (see Figure 26).

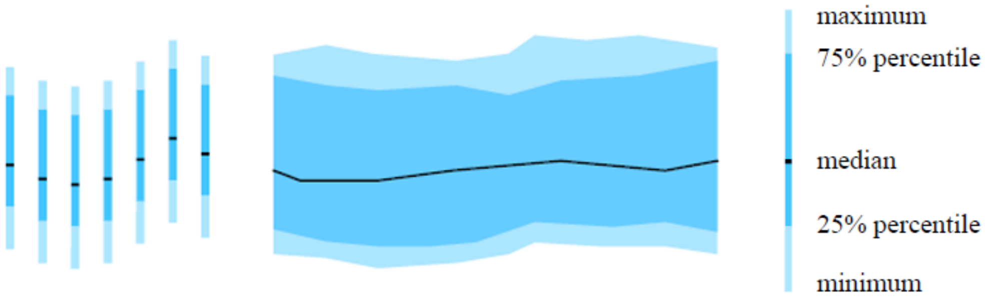

Figure 26: Redesign of the Tukey box-plot - inspired by [Tufte, 1983]. Big amounts of these small box-plot-bars can be connected to a full data stream representation of high-frequency data on very limited space [Bade et al., 2004].

Midgaard can also be used to visualize treatment plans. Figure 27 shows a treatment plan and its sub-plan. Plan characteristics such as preferences, conditions, intentions, and effects are given in tabular form.

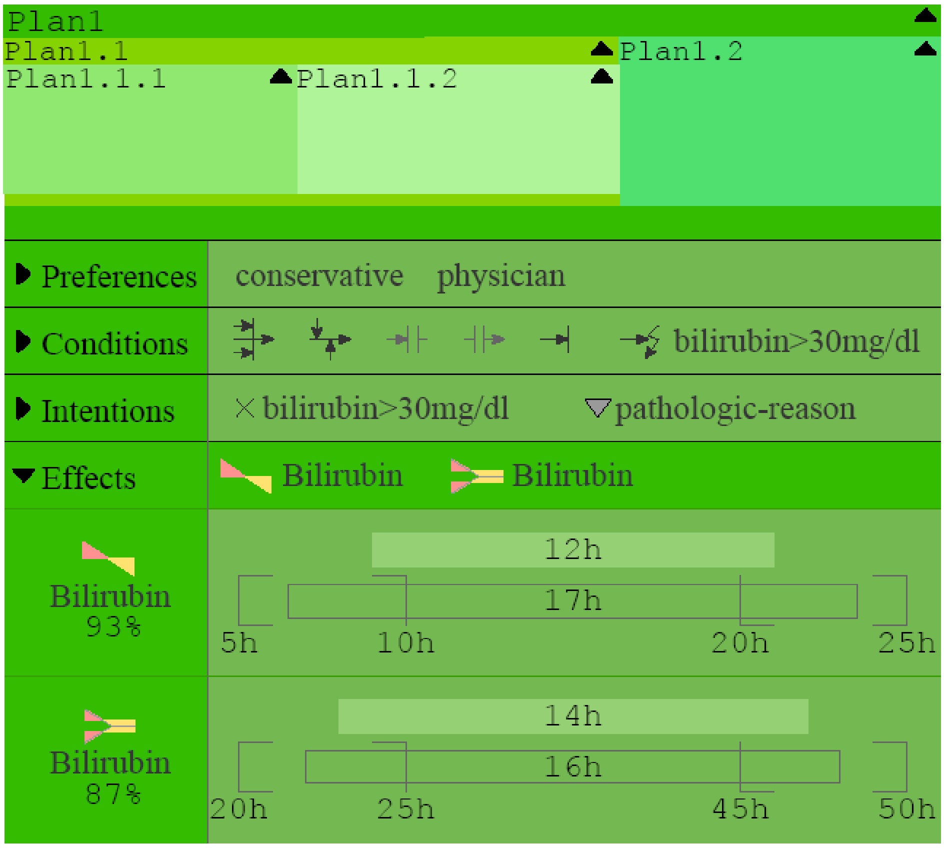

Figure 27: Visualization of a treatment plan as well as relevant plan characteristics (e.g. preferences, conditions, intentions, and effects) and the hierarchical structure of sub-plans. The estimated time of attainment of desired effects are visualized by special time-glyphs [Bade, 2002].

Moreover, the Midgaard prototype (see Figure 28) provides representations of measurement deviation, trust-ability of data points, and missing data as well as a number of interactions, like browsing, Pan+Zoom in time, Focus+Context, and integrated semantic zoom functionality (representations with different levels of detail). In addition, Midgaard provides different views for physicians, nursing staff, and patients.

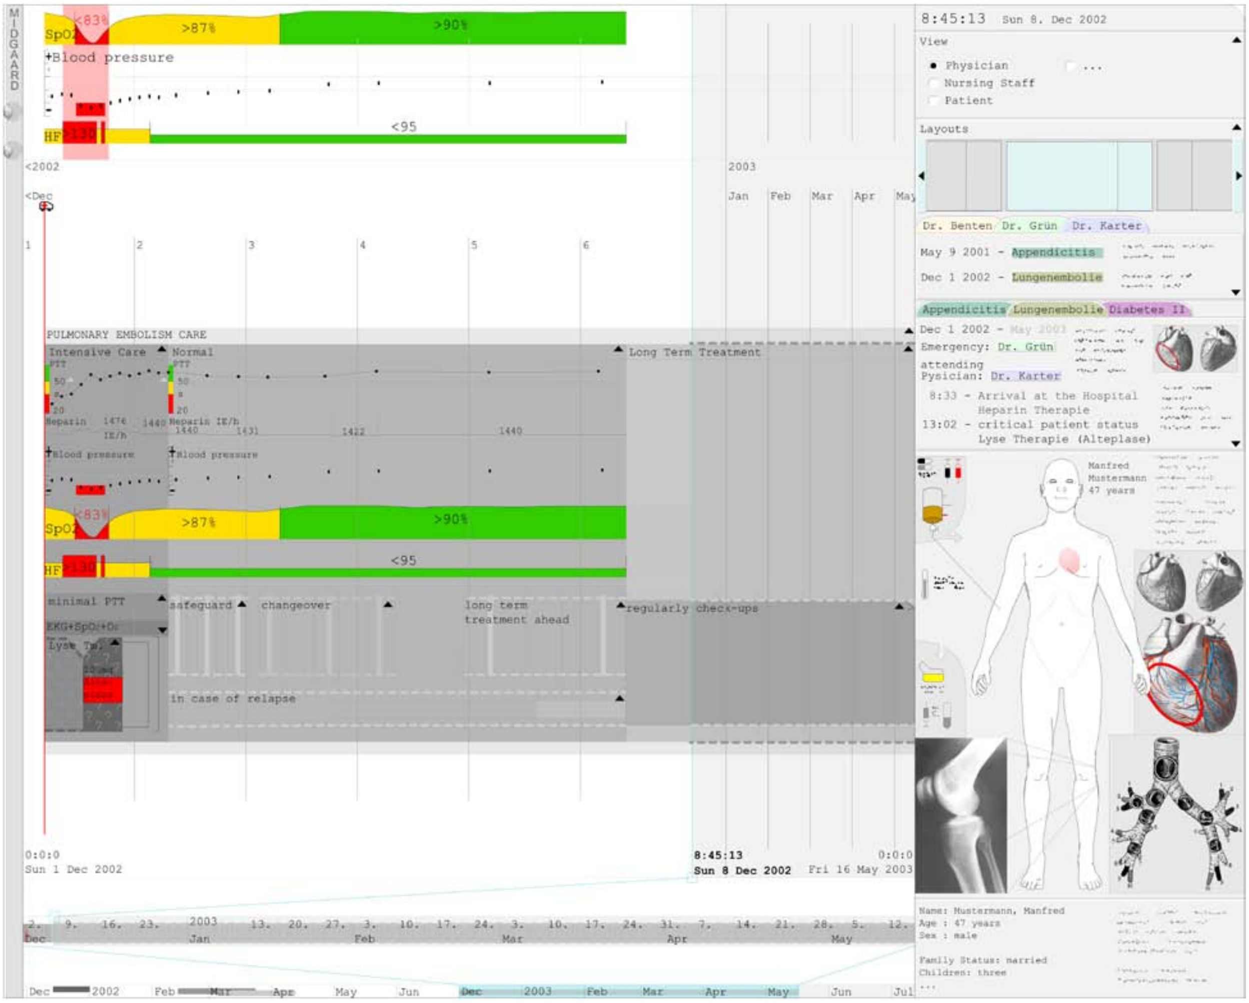

Figure 28: Screenshot of the Midgaard prototype [Bade et al., 2004].

5.3.3. CareVis

The CareVis interface [Aigner & Miksch, 2006] consists of three Views: the *Logical View* and the *Temporal View* along with a *Quick View* panel (see Figure 29).

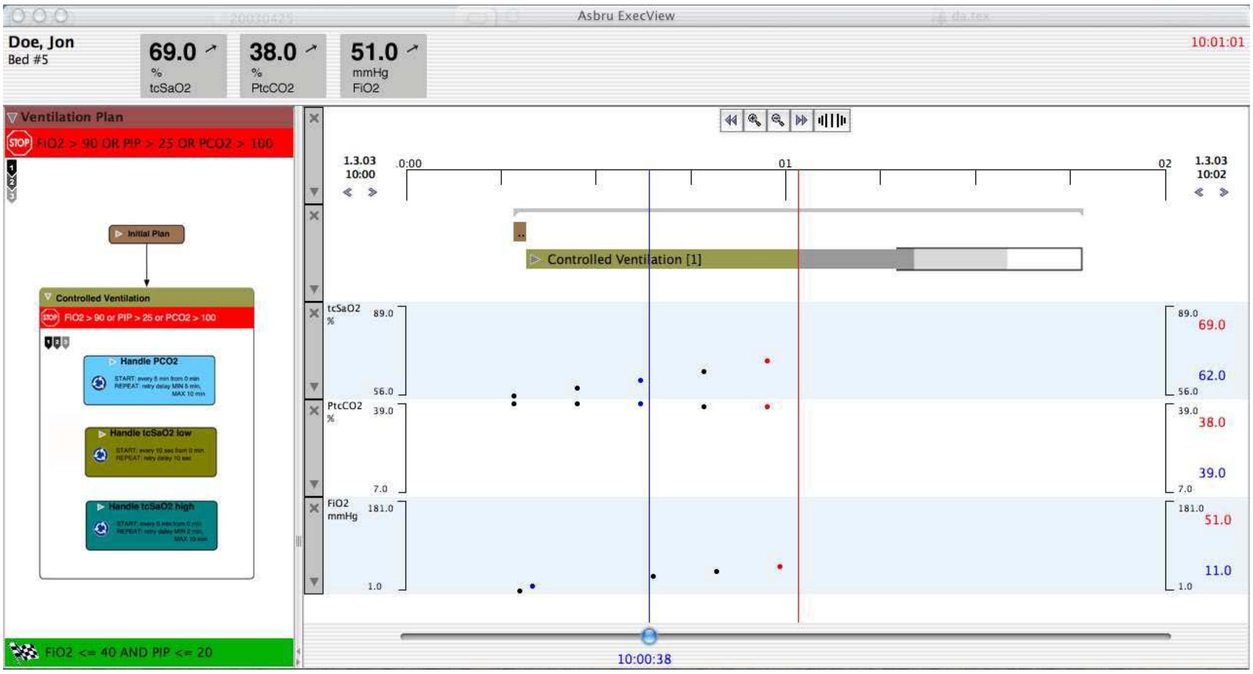

Figure 29: CareVis application window showing the execution of a plan [Aigner & Miksch, 2006].

The **Logical View** is on the left-hand side of the interface. It provides a representation of the treatment plan specification data which is called AsbruFlow (see Figure 30). AsbruFlow is based on flow-chart-like Clinical Algorithm Maps [Hadorn, 1995] extended by Focus+Context techniques to avoid an overcrowded appearance. Moreover, elaborated symbols indicate the execution order of plans, i.e. a set of six visual elements represent the different characteristics of Asbru plans. An Asbru plan is represented by a rounded rectangle filled with the color of the plan. Cyclical plans are indicated by an additional roundabout icon and the repeat specification in textual form within the rectangle. A physician icon indicates a user-performed plan and an ask step is marked by a question mark together with the text "Ask" followed by the parameter to be entered into the system. If-Then-Else constructs are visualized as hexagons and the textual condition. The visualization of conditions and the execution order is accomplished by enclosing frames: the topmost bar filled with the plan color contains the title of the plan. Below the title the abort condition is written which is depicted by a red bar with a stop sign icon on the left side. A green bar at the bottom of the plan represents the complete condition which is visualized by a checked finish flag icon on the left. The largest part of the enclosing frame represents the plan body along with the Execution Sequence Indicator. There are four possible symbols indicating the execution order. Moreover, AsbruFlow provides several interactive features such as to collapse and enlarge sections, Focus+Context techniques etc.

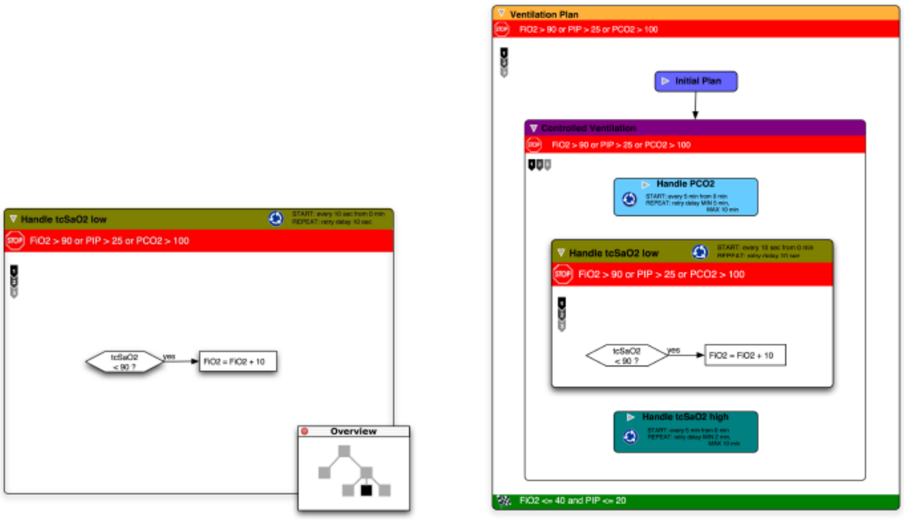
Handle tcSaO2 low
START: every 10 sec from 0 min
REPEAT: retry delay 10 sec
FiO2 > 90 or PIP > 25 or PCO2 > 100
tcSaO2 < 90 ?
yes
FiO2 = FiO2 + 10
Overview
Ventilation Plan
FiO2 > 90 or PIP > 25 or PCO2 > 100
Initial Plan
Controlled Ventilation
FiO2 > 90 or PIP > 25 or PCO2 > 100
Handle PCO2
START: every 5 min from 0 min
REPEAT: retry delay MIN 5 min, MAX 10 min
Handle tcSaO2 low
START: every 10 sec from 0 min
REPEAT: retry delay 10 sec
FiO2 > 90 or PIP > 25 or PCO2 > 100
tcSaO2 < 90 ?
yes
FiO2 = FiO2 + 10
Handle tcSaO2 high
START: every 5 min from 0 min
REPEAT: retry delay MIN 2 min, MAX 10 min
FiO2 <= 40 and PIP <= 20

Figure 30: The Logical View (AsbruFlow) showing parts of the Asbru plan for artificial ventilation of newborn infants – Overview+Detail mode (left) vs. Fisheye view mode (right) [Aigner & Miksch, 2006].

The **Temporal View** covers the representation of both, time oriented aspects of Asbru plans as well as of parameters and variables over time. The visualized interval can span from a point of time in the past to a point of time in the future. Thus, only plans with time annotations can be visualized accurately in the future, but no patient data. The Temporal View consists of a time scale, a facet that contains the treatment plan elements, and diverse facets containing different plan parameters and variables measured or computed over time. The time scale provides a fisheye feature, which magnifies a part of the scale interval in order to emphasize an area of interest, which allows for a detailed examination while preserving the context information.

The temporal representation of treatment plans is derived from the visualization of LifeLines [Plaisant et al., 1998] providing additional features such as displaying the hierarchical decomposition and Asbru's complex time annotations. The expandable representation of a treatment plan to show its sub-plans is called LifeLines+. In addition, the visualization of the temporal uncertainties future plans is called PlanningLines. A PlanningLines glyph allows for depicting Asbru's time annotations like ESS, LSS, EFS, LFS, MaxDur, and MinDur. It consists of two encapsulated bars representing the minimum and maximum duration, which are bounded by two caps that represent the start and end intervals (see Figure 31).

The visualization of the patient data is accomplished by facets below the treatment plan elements showing the course of some parameters over time. In addition, the **Quick View** panel in the top of the application window allows for monitoring the most important values at a prominent position. It displays currently valid variable and parameter values. The numerical value is displayed along with its name, its unit, and a trend indicator.

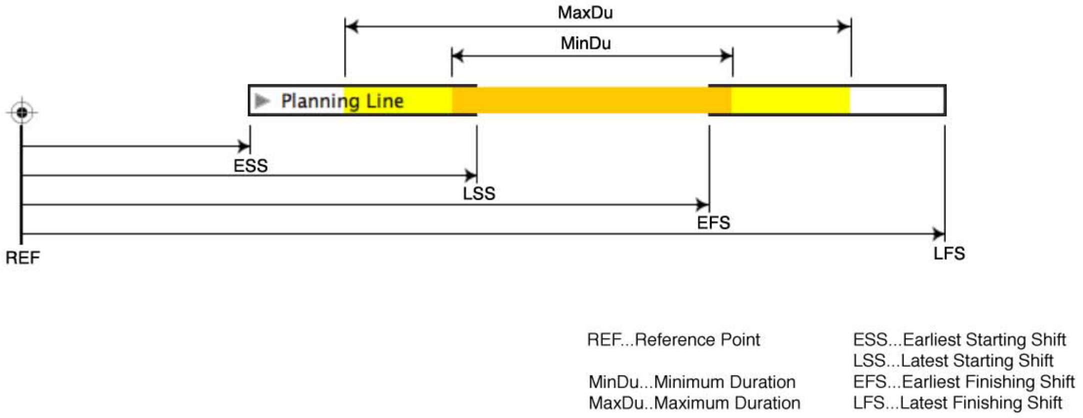
MaxDu
MinDu
Planning Line
ESS
LSS
EFS
REF
LFS
REF...Reference Point
MinDu...Minimum Duration
MaxDu..Maximum Duration
ESS...Earliest Starting Shift
LSS...Latest Starting Shift
EFS...Earliest Finishing Shift
LFS...Latest Finishing Shift

Figure 31: The PlanningLines glyph depicts the earliest starting shift, the latest starting shift, the earliest finishing shift, the latest finishing shift, the maximum duration, and the minimum duration of a treatment plan [Aigner & Miksch, 2006].

5.4. Microsoft Health Common User Interface

The Microsoft Health Common User Interface provides design guidance and toolkit controls with the objective to standardize display elements of health applications and thus, ensuring patient safety. For instance, standardized display elements (e.g. patient identification information) allow for easily switching between different applications and thus, increase clinical effectiveness and patient safety.

In collaboration with the National Health Services [NHS] in England (in context of the NHS Common User Interface (CUI) Programme) they formulated a set of patient safety principles [Microsoft Corporation]:

1. The correct identification of a patient and a matching of the patient to their care elements
2. Prevention of patient care hand-over errors and safety during transition of care
3. Assuring medication accuracy during the giving of care to a patient
4. Performance of correct procedure at correct body site

With respect to these principles they have developed a set of design guidelines and Toolkit controls.

Microsoft Health CUI Design Guidance is a set of user interface guidelines and recommendations with can be applied to new software applications or used to evaluate clinical safety of existing software applications. For instance, guidelines on the correct identification of patients recommend to displaying the patient's Family name in all uppercase letters to clearly distinguish it from the Given name, etc. (see Figure 32).

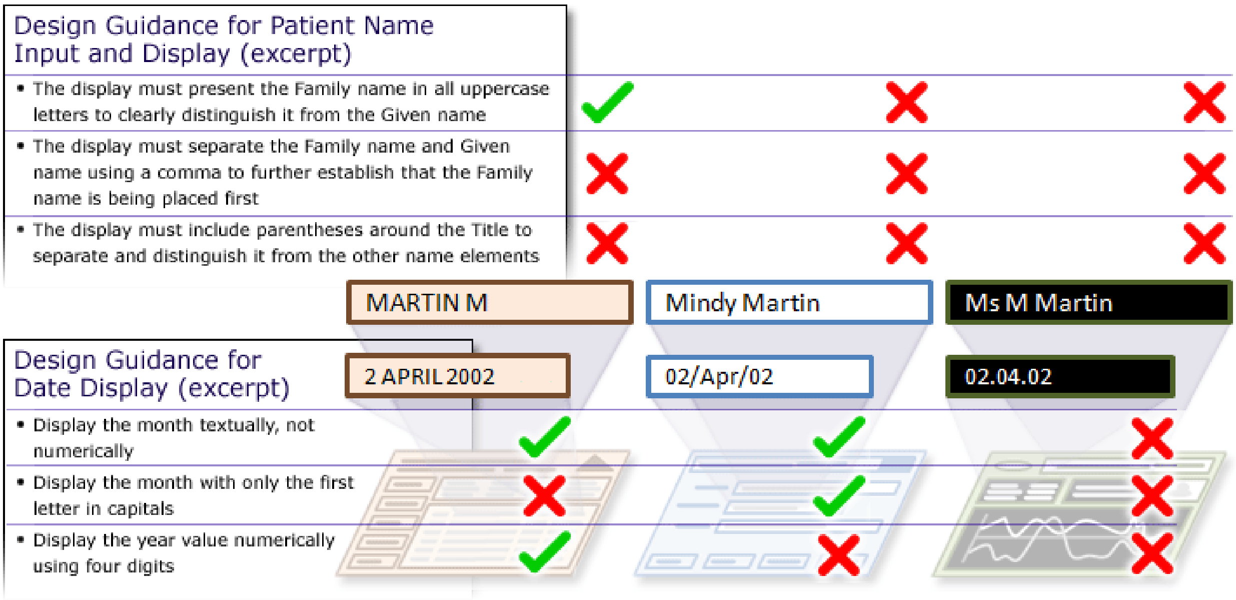

Design Guidance for Patient Name Input and Display (excerpt)	MARTIN M	Mindy Martin	Ms M Martin
The display must present the Family name in all uppercase letters to clearly distinguish it from the Given name	✓	✗	✗
The display must separate the Family name and Given name using a comma to further establish that the Family name is being placed first	✗	✗	✗
The display must include parentheses around the Title to separate and distinguish it from the other name elements	✗	✗	✗

Design Guidance for Date Display (excerpt)	2 APRIL 2002	02/Apr/02	02.04.02
Display the month textually, not numerically	✓	✓	✗
Display the month with only the first letter in capitals	✗	✓	✗
Display the year value numerically using four digits	✓	✗	✗

Figure 32: Microsoft Health Common User Interface Design Guidance for patient name and date display [Microsoft Corporation].

Microsoft offers detailed guidelines on:

1. Clinical noting and terminology (including guidance on displaying graphs and tables)
2. Consistent navigation
3. Medications
4. Patient Identification
5. Miscellaneous (including accessibility principles and accessibility checklist)

For instance, guidelines on displaying graphs and tables focus on [Microsoft Corporation]:

- Providing appropriate layout and formatting of graphs and tables for data values and data series
- Providing consistent scaling and appropriate axes
- Analyzing data values against normal ranges, and in a series of data values taken over time, including the use of interpolation
- Representing the same data values within both graphs and tables
- Demonstrating where access to supporting contextual data will be necessary

Microsoft Health CUI Toolkit Controls are ready-made user interface controls which fully adopt the design guidelines, for instance a “NameLabel” control or a “DateLabel” control. These controls can be customized and used for the development of new health application user interfaces.

5.5. Visualization of Hierarchical Data

As an Asbru plan can be decomposed into sub-plans, which in turn contain sub-plans themselves, the plan and its sub-plans build a hierarchical structure. In the following subsection we outline existing techniques to visualize hierarchical data.

5.5.1. Trees and Treemaps

A tree is a visualization technique which represents the hierarchy by a directed and acyclic graph with one root node (see Figure 33(b)). Each node of the graph (except for the root node) has exactly one incoming edge from its parent node. Thus, a node B is called child of node A if there is an edge from node A to node B. In turn, node A is called parent of node B. Each node may have 0-many child nodes and those without child nodes are called leaf.
Usually these graphs are aligned from top to bottom (starting with the roof node) representing each level of the hierarchy at a separate row.
Other approaches such as hyperbolic trees [Lamping et al., 1995] place the root node in the center and align child nodes on circular layers around the roof node (see Figure 34). Thus, the number of nodes grows exponentially with each hierarchy level which causes a reduction of the node size as moving outwards. To this end, hyperbolic trees use focus and context techniques which allow the user to change the focus and explore each region of the graph in detail while providing less detailed context information.

A treemap [Shneiderman, 1992] is a rectangular, space-filling approach for visualizing hierarchical data (see Figure 33(c) and Figure 35). It uses 2D space to lay out the area of the root element which encapsulates the child elements. The size of each node is determined in relation to all other nodes; it represents a given attribute of the node (e.g., file size in case of the hierarchy of directories on a hard drive). Another way to encode additional information with treemaps is the use of color. Thus, treemaps are an effective way to visualize different attributes of nodes and enable the user to compare nodes and sub-trees on different levels of the hierarchy.

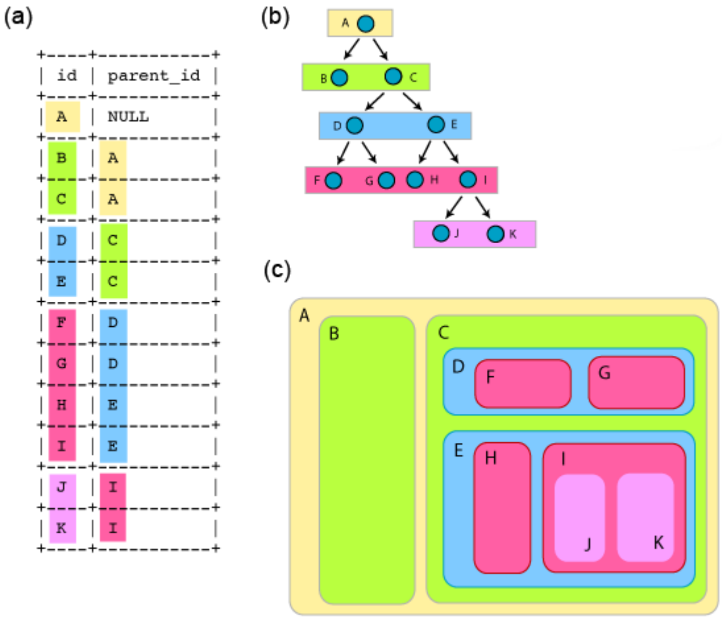

id	parent_id
A	NULL
B	A
C	A
D	C
E	C
F	D
G	D
H	E
I	E
J	I
K	I

Figure 33: Most common ways to visualize hierarchical data: (a) listing all elements together with their parent elements, (b) a tree-layout starting with the root element A from which edges lead to the child elements B and C, and (c) a treemap which uses 2D space to encapsulate child nodes into the area of their parent node [Orsini, 2006].

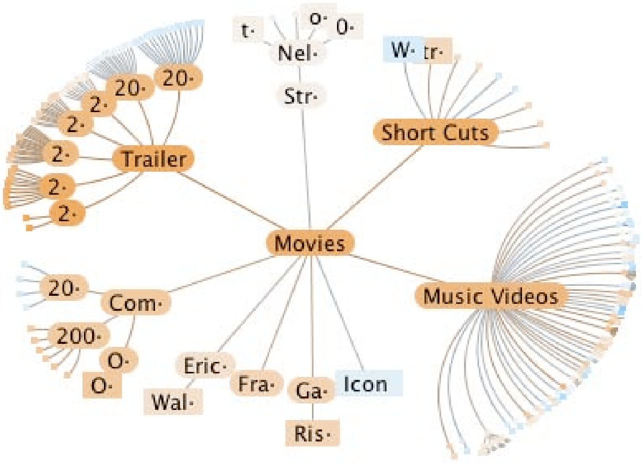

Figure 34: Hyperbolic Tree: the root node is located in the center and child nodes are aligned in circular layers around the root node. Interactive focus and context techniques allow for navigating the graph and investigating each node/area in detail while providing less detailed context information [Randelshofer, 2009].

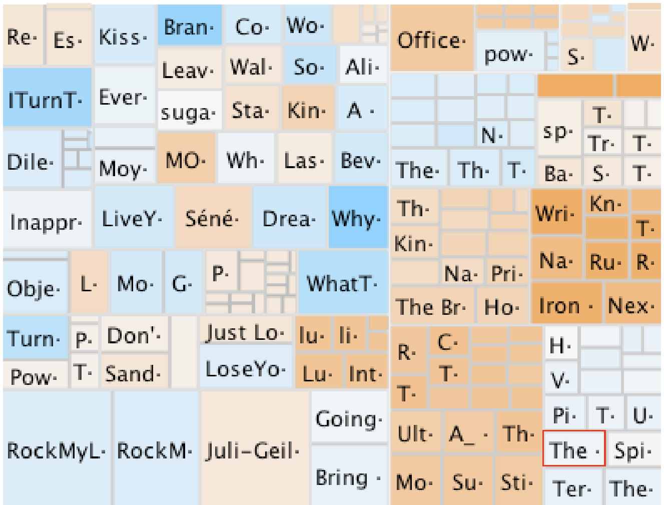

Figure 35: Treemap: a rectangular, space-filling visualization of hierarchical data. The size and color of the nodes encode additional information [Randelshofer, 2009].

5.5.2. Timeline Trees

Timeline Trees [Burch et al., 2008] address the challenge to visualize hierarchical structures over time (see Figure 36). It was designed to visualize market baskets of one or more customers as a sequence of transactions that occur over time. The hierarchy is represented by a horizontally aligned tree structure with interactively selectable nodes. By clicking on an intermediate node the sub-tree starting at this node is collapsed and in turn, a node that is currently displayed as a leaf is expanded when clicking on it. This selection of a hierarchy level is propagated to all other views which allows the user to explore the different levels of the hierarchy. The sequence of transactions is visualized as an extension at the right side of the tree structure where transactions are laid out along a timeline. The transactions are represented by rectangles with different size and color encoding given attributes. Between the tree and the timeline view there are thumbnail representations of the timeline view for each leaf or collapsed node displaying only those transactions that include this specific leaf or collapsed node.

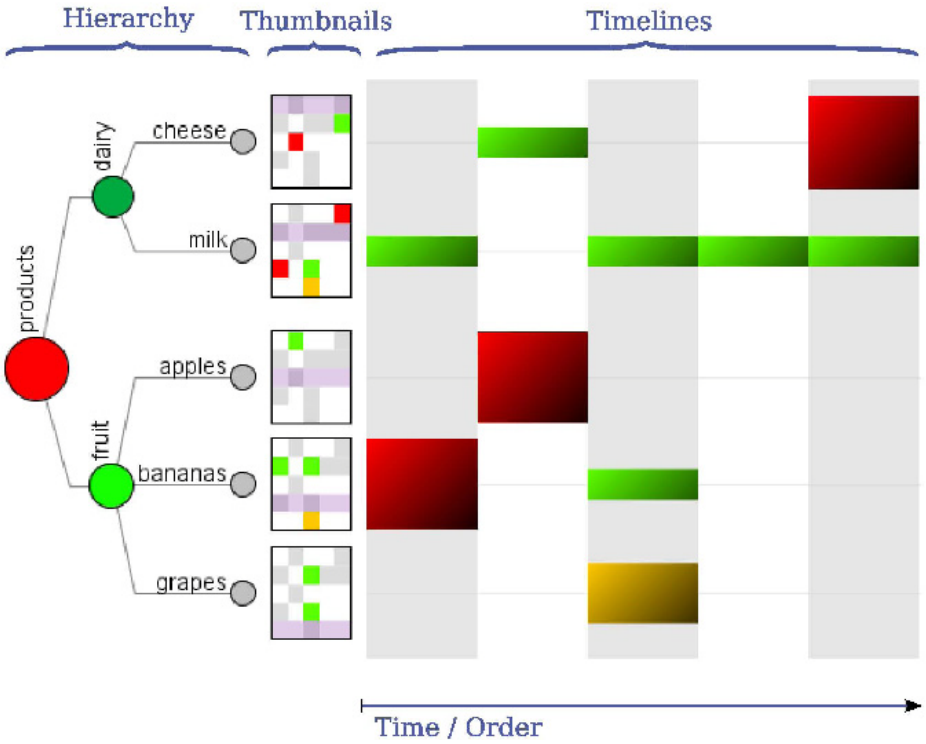

Figure 36: Timeline Tree: the hierarchy visualization on the left hand side is interactive and allows for expanding and collapsing nodes to select different layers of the hierarchy. The selection is propagated to the two other views, the thumbnails and the timelines [Burch et al., 2008].

6. Shortcomings of Existing Approaches

In the previous sections we presented 18 research approaches of visualizing treatment plans, medical records, or both, patient data in combination with applied treatment plans. Next, we will outline how well the introduced visualization techniques match the given list of requirements. See Table 1 for an overview.

The first seven visualization approaches focus on the visualization of treatment plans. *Clinical Algorithm Maps*, the *Tallis* toolset, *Protégé*, *GUIDE*, and *Glare* offer a flowchart-like representation of clinical algorithms suited to visualize the logics and execution sequence of plans, but neither the hierarchy of treatment plans nor complex plan characteristics. The purpose of *VisiGuide* is to browse and present a large amount of guidelines and allow the user to explore a single guideline in detail. Thus, it displays the guideline content and different characteristics like conditions and intentions in textual form. However, there is no actual visualization of guideline plans. Furthermore, *AsbruView* is able to visualize all relevant plan characteristics, but does not visualize patient data.

On the other hand, there is a variety of approaches visualizing patient-related information. The *Graphical Summary of Patient Status* draws a comprehensive picture of a patient’s condition over time, but is very space consuming; *LifeLines* visualizes only some qualitative parameters of patient records over time, and *PatternFinder* is focused on filtering the patient data and visualizing only a small set of special events. *KNAVE* represents different patient parameters over time and allows the use to explore them; it is not suited to give a comprehensive overview. *VISITOR* is aimed at visualizing parameters of multiple patient records, which does not serve our requirements, either. Furthermore, *VIE-VISU* is designed for the very special purpose of visualizing neonatal intensive care data. It uses glyphs to represent the patient's condition; unfortunately, the correct interpretation of these glyphs demands for a huge learning effort. The 3D visualization technique of the *Interactive Parallel Bar Charts* approach allows for exploring clinical parameters of hemodialysis sessions. This technique is quite space consuming; besides, problems of 3D representations (e.g. occlusions) are still not solved satisfactorily. *Gravi++* visualizes the change of the patient over time by using animation and traces but it is not suited to give a comprehensive overview of the patient's condition.

There are only a few tools visualizing patient data in combination with applied treatment plans.
The *Guideline Overview Tool* fails in visualizing plan hierarchies, complex plan characteristics, as well as a comprehensive picture of the patient's condition.

Midgaard visualizes detailed and compact patient data over time as well as the hierarchical structure of plans, complex characteristics (e.g., plan states and conditions), and execution sequence. Still, the representation of treatment plans requires some learning effort and there are no means to explore the interrelations of applied treatment plans and the patient's condition. Finally, *CareVis* uses sophisticated visualization techniques to represent relevant characteristics of treatment plans but fails in drawing a comprehensive picture of a patient's condition.

The *Microsoft Health Common User Interface* provides useful guidelines on standardized data labels and the representation of very basic graphical user interface components, such as graphs and tables. It is understood that these guidelines can only cover common user interface elements, but not complex information visualization techniques as required for the representation of treatment plans in combination with patient data as described in this paper. However, we will certainly respect standards of user interface design and consistent representation of common interface components.

Table 1. Comparison of existing approaches and requirements.

	Hierarchy of plans	Execution sequence	Plan characteristics	Comprehensive overview	Patient's course over time	Space saving visualization	Interrelations
Clinical Algorithm Maps		●					
Protégé		●	○				
Tallis Toolset		●					
GUIDE		●					
GLARE		●					
VisiGuide			○				
AsbruView	●	●	●				
Graphical Summary of Patient Status				●	●		
LifeLines					●	●	
PatternFinder						●	
KNAVE				○	●		
VISITORS					●		
VIE-VISU				○	○		
Interactive Parallel Bar Charts					●		
Gravi++					●		
Guideline Overview Tool		○			●		
Midgaard	○	○	○	●	●	●	
CareVis	●	●	●		●		

7. Summary of Requirements

In order to meet the above mentioned requirements, sophisticated techniques have to be applied to provide:

- An information rich overview of a patient's condition. It is crucial to provide a comprehensive and well-arranged representation of important patients' parameters to allow for an efficient investigation of the patient's condition.
- Changes of a patient's condition over time. Changes of parameters have to be illustrated to allow for a detailed investigation of improvement or worsening of the patient's condition. This is of particular importance for the time span during and after applied treatment plans.
- A comprehensive visualization of relevant characteristics of applied treatment plans. Treatment plans are complex constructs which embody diverse characteristics relevant for execution, for instance execution sequence and different conditions and effects. A compact and intelligible visualization of the treatment plans together with these characteristics has to be found.

- Insights into relations between them. Providing means to investigate relations between treatment plans and a patient's condition, for instance to evaluate the success of different treatment plans, is a challenging issue which was not addressed yet.

Moreover, all these data and information has to be visualized in combination on very limited display space.
Selected visualization approaches cover certain aspects very well, for instance, *CareVis* provides solutions for visualizing complex information about applied treatment plans in an intelligible and comprehensive way. We took this approach as a starting point to develop a consistent visualization. However, most challenging is the task of making interrelations between a patient's condition and applied treatment plans transparent. Besides an appropriate visualization opposing patient data to treatment plans, we provide elaborated interactive means to allow the user for exploring the data and in doing so, enable him/her to reveal new insights and to formulate hypotheses about how applied treatment plans effect the patient's condition or to compare different therapies.

8. The Visualization Module Supplied by TUW

In order to provide a tool that supports every aspect of treatment plan execution, we decided to extend the possibilities of CareVis [Aigner & Miksch, 2006] which covers the nesting of treatment plans, the execution order, relevant plan characteristics, and temporal constraints. For this reason we have developed a conceptual extension of CareVis, which, on the one hand, provides means to explore and assess success or failure of applied treatment plans and effects of clinical actions on a patient's condition, and, on the other hand, allows for the investigation and comparison of two or more patients at the same time.

Extending the possibilities of CareVis resulted in adding value to an already very powerful tool, and thus, to provide a tool to support every single aspect important in the application of computer-executable treatment plans. On the other hand, it was a great challenge to find ways to smoothly integrate the visual encoding of additional information into the user interface of CareVis. CareVis already provides three different views to communicate specific information, namely the logical view to visualize the logics of treatment plans, a view to show the hierarchical structure of these plans, and the temporal view to represent complex time constraints of treatment plans (see Section 5.3.3).

Our extensions serve a two-fold purpose: (1) to provide visual and interactive means for assessing effects of applied treatment plans on patients using color-coded distance information and slope and (2) to compare multiple patients at the same time (see Figure 37 for the composition of the GUI of the visualization module). In the following subsections we give detailed description of the visualization module.

8.1. User Interface Composition

In this section, we use examples from the domain of artificial ventilation to demonstrate our approach, because the ReMINE protocols and simulated data matching these protocols are not sufficient to demonstrate all the features of the visualization (see Table 2). Section 9 demonstrates the application of our approach on the ReMINE protocol models with simulated data.

We designed the user interface with respect to the findings of the analysis of data characteristics and requirements. To this end, the user interface (see Figure 37) provides:

- Charts of the patient's parameter values (see Section 8.2.1),
- A graphical representation of treatment plans:
    - Currently applied treatment plans,
    - Treatment plans that have been applied in the past,
    - Treatment plans which are to be applied in the near future,
    - Plans are represented with all relevant characteristics:
        - Execution order and plan nesting (see Section 8.2.4),
        - Temporal constraints (see Section 8.2.5),
        - Clinical actions and the time of execution (see Section 8.2.6), and
        - Success or failure of the plan (see Section 8.2.7),
- Patient parameters can dynamically be added and removed from the visualization,
- The representation of the parameter charts can be manipulated, which allows for the visualization of actual parameter values at the time of measurement, of a connected line chart to enhance the visibility of trends, and of a combination of both,
- Selected clinical actions can be highlighted in the treatment plans to give an overview as to when and in which context these actions have been applied, which allows the user to compare the effects of these events, and
- Interactive means are provided to compare selected treatment plans as well as selected actions in detail (see Section 8.3).

Table 2. Comparison of the features of the interactive visualization and which protocols and the corresponding data are suited to demonstrate them.

	Hierarchical nesting of plans	Logical structure of plans	Intentions of plans	Execution order of plans	Color-coded distance to intended value	Color-coded progress from initial value	Color-coded slope	Aligning plans or actions	Investigation of effects of clinical actions	Detection of color patterns with the focus window
ReMINE BirthAsbru protocol	●	●		●						
ReMINE InfectionAsbru protocol	●	●		●						
ReMINE MedicalAsbru protocol	●	●		●						
ReMINE StrokeAsbru protocol	●	●		●						
Example Ventilation protocol	●	●	●	●	●	●	●	●	●	●

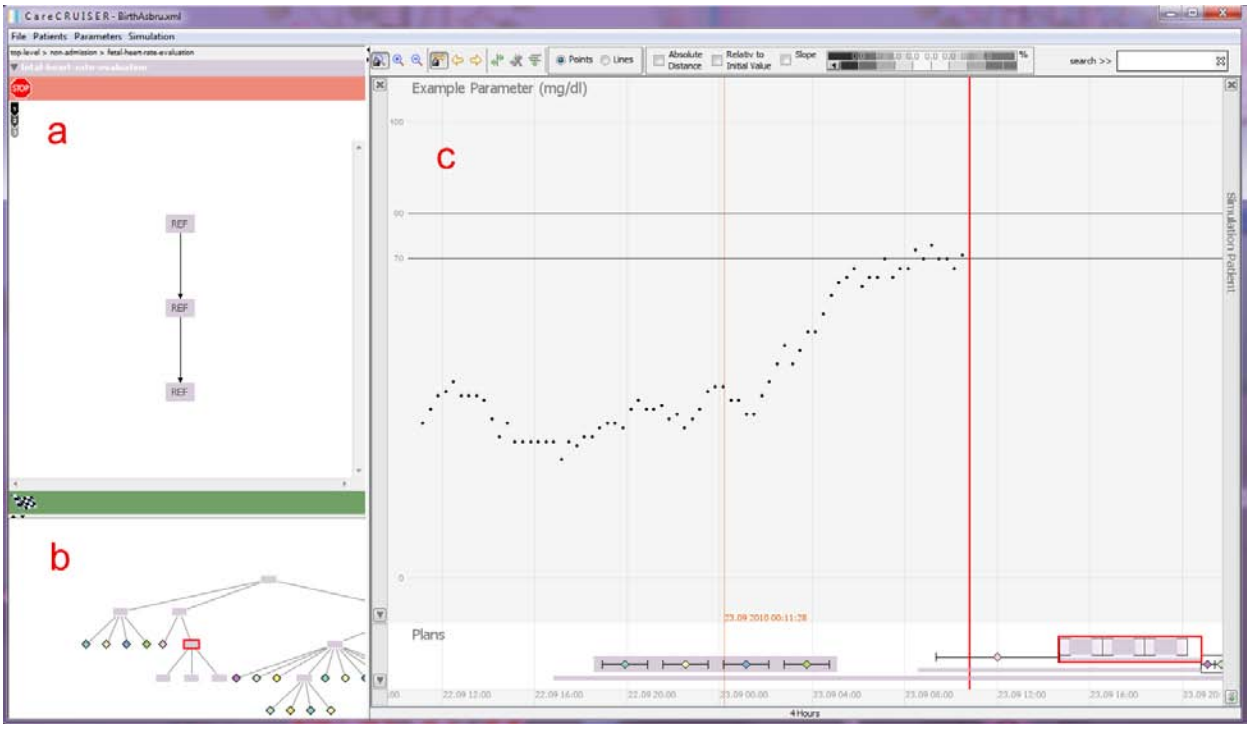

Figure 37: Composition of the user interface of the visualization module. The logical view (a) communicates the logical structure of treatment plan execution by means of a flowchart-like representation as well as diverse plan characteristics (such as conditions, intentions, and execution order of sub-plans). This view provides conditions and intentions of the selected plan in textual form (i.e., the filter-precondition next to the red stop-sign, complete-condition next to finishing-flag within the green bar, and intentions below the red bar next to the execution sequence indicator in the top-left corner). However, there are no conditions or intentions specified for this treatment plan. The lower left part (b) displays a tree graph to visualize the hierarchical structure of treatment plans and sub-plans [Aigner & Miksch, 2006]; the temporal view (c) focuses on the temporal-qualities of applied treatment plans and patient parameters. The thick red line separates the past from the future. In the temporal view we provide step-wise interactive means to explore the effects of applied treatment plans on the patient's condition.

8.2. Visual Encoding

Information Visualization (InfoVis) uses visual attributes such as position, size, color, orientation, etc. to communicate complex ideas precisely [Few, 2004]. Thus, the superior goal is to design a visualization that encodes all the required information and data in an easily understandable and consistent way. Visual artefacts which do not add relevant information to the visualization obfuscate the information and are indications of bad style. Therefore, each visual attribute we use within the visualization aims at communicating relevant information as described in the following subsections.

8.2.1. Patient's Parameters

We visualize the course of patient parameter values as simple charts over time which is a well known and intuitive representation. Thus, the representation does not require any extra effort to get familiar with. Besides, it provides raw data without modifications or abstractions which is important for physicians as they are used to assess a patient's condition by these raw values.

Displaying the parameters affected by an applied treatment plan (i.e., parameters that are subject of the plan’s intentions) is mandatory to allow for a detailed investigation of the patient’s condition and the effects of the applied treatment plan. The selection of other relevant patient parameters, however, cannot be generalized. Therefore, we provide a possibility for physicians to select and inspect parameters of interest, and thus, to tailor the visualization to their specific needs.

8.2.2. Hierarchical Nesting of Treatment Plans

To represent the hierarchical structure of a treatment plan and its sub-plans we provide a tree graph (see Figure 38) in the hierarchical view. Tree graphs are a well-known and appropriate visualization technique to communicate hierarchical structures. To facilitate the coordination between views, selected plan elements are highlighted synchronously in each view.

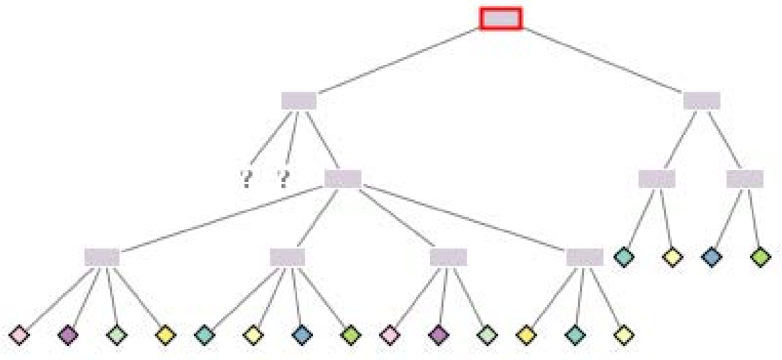

Figure 38: Tree graph of the treatment plan and its sub-plans. Plans are represented by rectangles, actions of assessing information about the patient (e.g., anamnesis) are represented by question marks, and clinical actions (e.g., administration of a certain drug) are represented by diamonds.

8.2.3. Logical Structure of Treatment Plans

The logical structure of a given treatment plan is visualized in the logical view (see Figure 39). The name of the selected plan is displayed in the header below a breadcrumbs trail to facilitate orientation within the hierarchical plan structure. Below the abort condition of the plan is given in textual form and indicated by the red stop-sign. At the top of the main window there is the execution sequence indicator (see Figure 41) and a textual description of the intentions specified for this treatment plan. The complete condition is represented by a checked finish flag icon in the green bar at the bottom of the plan and also given in textual form (in Figure 39 no complete condition is specified for the given plan).

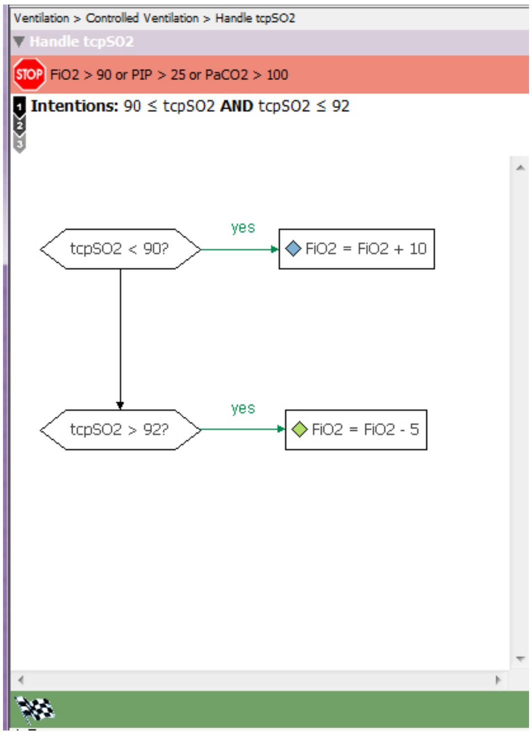

Ventilation > Controlled Ventilation > Handle tcpSO2

Handle tcpSO2

STOP FiO2 > 90 or PIP > 25 or PaCO2 > 100

Intentions: 90 ≤ tcpSO2 **AND** tcpSO2 ≤ 92

Figure 39: The logical view provides the logic of treatment plan execution in a flowchart-like representation.

8.2.4. Execution Order of Treatment Plans

With respect to the relevant temporal dimension, we lay out treatment plans along a timeline (the x-axis). This is a well-established and intuitive way of visualizing time. Other visualization options are, for instance, animation and traces, both of which make the comparison of two points in time more difficult and fail in giving an overview of a time span in one picture. Other ways of visualizing time are small multiples or multiple views which can display only a very limited number of points in time. Thus, these methods do not meet our requirements.

Treatment plans are represented by rectangles. This visual encoding allows for the representation of the exact start- and end-time of the plan on the horizontal timeline. The following subsections contain an outline of the visual encoding of different types of plans.

8.2.4.1. Treatment Plans with Strictly Predefined Execution Order

Asbru Plans are complex constructs with nested sub-plans; each type of plan determines a specific order of executing these sub-plans. In the case of sequential plans and parallel plans, the execution order of their sub-plans is strictly predefined.

Sequential Plans

All sub-plans of a sequential plan are to be executed one after the other in a predefined order. The sequential execution is visually encoded by laying out one sub-plan after the other along the timeline. In case of an expanded parent-plan, the sub-plans are aligned above the parent-plan; when selecting the parent-plan, a red rectangle borders the plan together with its sub-plans (see Figure 40).

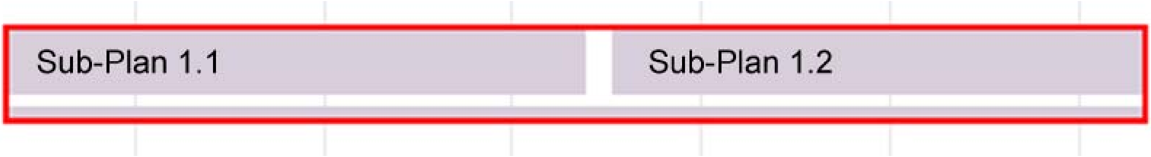

Figure 40: Sequential plan: the sub-plans are laid out on a timeline (x-axis) one after the other which shows that the plans are carried out in sequence.

In addition, the execution sequence is indicated by different symbols in the upper left corner of the logical view (see Figure 37). There are four possible symbols indicating the execution order (see Figure 41).

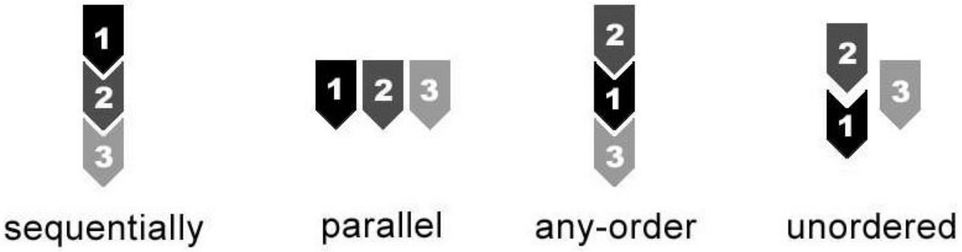

Figure 41: Execution sequence indicator: four different symbols indicating four different execution sequences [Aigner & Miksch, 2006].

Parallel Plans

All sub-plans of a parallel plan get initialized at the beginning and are processed in parallel. In contrast to sequential plans, the sub-plans of parallel plans are aligned above each other indicating the parallel execution (see Figure 42).

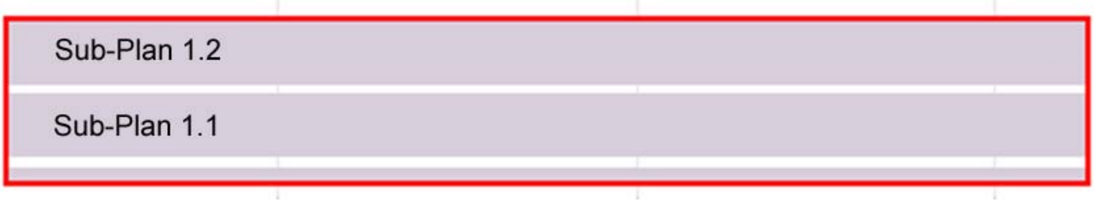

Figure 42: Parallel plan: The sub-plans are stacked vertically which is to indicate that the plans are carried out in parallel.

Cyclical Plans

A cyclical plan contains a sub-plan that can be repeated. The time interval for which the plan is repeated, the start time, the end time, and the duration of the particular plan instance in each cycle are predetermined. The repetition of the sub-plan is similar to the execution of sub-plans of a sequential plan which leads to a similar visual encoding (see Figure 43).

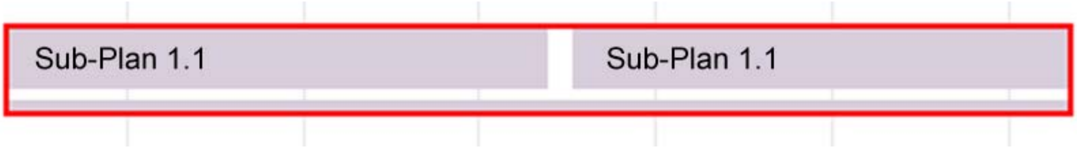

Figure 43: Cyclical plan: a cyclical plan contains a sub-plan that can be repeated. The visual encoding is similar to that of sequential plans but the plans that are executed in sequence are identical.

8.2.4.2. Treatment Plans with Flexible Execution Order

The execution order of any-order plans and unordered plans is not strictly predefined. All sub-plans of an any-order plan are processed one after the other, but the order is arbitrary. Sub-plans of an unordered plan can be executed in an arbitrary way – in sequence or parallel or some in sequence and others in parallel. The order is also arbitrary. They are visualized like sequential plans or parallel plans according to their actual execution order.

After defining the representation of plans and sub-plans we need to integrate relevant plan characteristics as outlined in the following subsections.

8.2.5. Temporal Constraints

The start and end time of plans is given implicitly by the bounds of the plan-rectangle and its position on the timeline. In contrast to common visualizations of hierarchical data, which often align child-elements as indented elements below their parent-elements, we went for a representation of sub-plans stacked below the parent-plan, which does not violate the time-axis. Other representations of hierarchical data – for instance, trees and treemaps (see Section 5.5.1) – cannot be used in this context because of their inability to visualize time. Timeline Trees (see Section 5.5.2) are a sophisticated method to visualize hierarchical data with respect to time, but add some complexity to navigation and fail in giving a complete overview of all hierarchy levels in one picture.

To visualize Asbru's complex time annotations of earliest starting shift, latest starting shift, earliest finishing shift, latest finishing shift, as well as the maximal and minimal duration of a plan, we have integrated the PlanningLine glyph (see Section 5.3.3) into our visualization (see Figure 44).

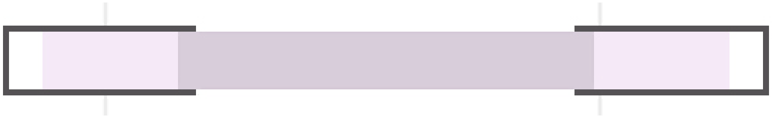

Figure 44: PlanningLine glyph [Aigner & Miksch, 2006]: the caps (the bracket-like forms at the far left and far right of the glyph) represent the earliest and latest starting shift and finishing shift; the maximum duration is indicated by the light gray area and the minimum duration is indicated by the dark gray area (see Figure 31).

8.2.6. Clinical Actions

For assessing the effects of clinical actions the user needs to know when, how often, and in which context these actions have been applied. To visually representation these actions we chose diamonds for two reasons (see Figure 45):

1. Clinical actions can take place at one single point in time. Thus, the visual representation must allow for the identification of this exact point in time; this is indicated by the peak of the diamond.

2. Moreover, diamonds have a body which makes them more visible than simple vertical lines (which is especially true for actions at the start and end of a treatment plan).

In addition, we extended the representation with respect to clinical actions that take more time to be carried out. The temporal bounds of these actions are indicated by whiskers that span the duration of the action and allow for the identification of the exact time of its start and end. Representing the start and end time by the width of the diamond would lead to line-thin diamonds for actions that take place at one single point in time, and thus, it would decrease the visibility of these actions. Similar to the representation of parallel plans, actions which are executed in parallel are aligned below each other. We assign a color (based on the qualitative color palettes proposed by Harrower and Brewer [Harrower & Brewer]) to each clinical action to facilitate their identification, for instance, two or more blue diamonds in a temporal treatment plan glyph indicate two or more instances of applying one clinical action (e.g., the application of a certain drug), in context of executing this treatment plan. Moreover, the color of clinical actions facilitates the coordination between the different views.

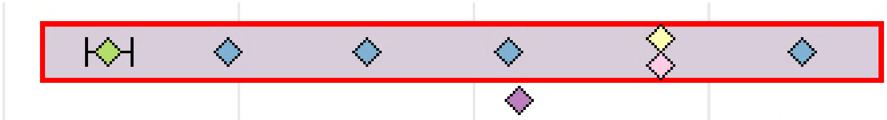

Figure 45: Clinical actions are represented by diamonds. In case of a clinical action that takes place at one single point in time, the peak of the diamond allows for the identification of the exact time of execution. On the other hand, the start and end time of a clinical action that takes longer is indicated by whiskers that span the duration of the action. Clinical actions that were applied to the patient but are not part of the treatment plan are laid out beneath the plan body.

Since the number of well-distinguishable colors is limited, we additionally support the identification of clinical actions by tooltip windows and the possibility to highlight all applications of a selected clinical action. However, the number of different colors is sufficient for our test datasets.

8.2.7. Effects of Applied Treatment Plans

In order to determine success or failure of an applied treatment plan we take advantage of the intentions specified for this plan (intentions represent goals and intended effects of the treatment plan). For instance, a treatment plan dealing with the controlled ventilation of new born children aims at stabilizing the patient's transcutaneous assessed oxygen saturation (tcpSO2) between 90% and 92% and the patient's pressure of carbon dioxide (PCO2) between 40 and 60 mmHg by adjusting the supply of inspired oxygen (FiO2) as well as adjusting the ventilation frequency (f). Moreover, the FiO2 is to be kept below 40% and the peak inspiratory pressure (PIP) below 60 cm H2O. This treatment plan can be decomposed into sub-plans of which each may aim at achieving selected aspects of the overall intentions, for instance, the sub-plan ‘Handle tcpSO2’ aims at stabilizing the patient’s tcpSO2 between 90% and 92% by adjusting the supply of FiO2. To visually encode how well an applied treatment plan succeeds in reaching the specified intentions and to make effects of applied clinical actions stand out more clearly, we use color-coding to highlight interesting events (i.e., distance information and slope) of the patient’s parameter development. We went for color-coded information highlighting (instead of, for instance, distance plots or plots of the first derivation) in order to maintain raw data charts without modifications or abstractions. This is important for physicians as they are used to assess a patient’s condition by these raw values. Besides the highlighting of color-coded information, the stepwise investigation of effects of applied treatment plans and clinical actions is supported by the following interactive means: filtering for interesting events of parameter values’ development, aligning plans or actions below each other to ease comparison, and selecting regions of interest by means of a focus window. In the following subsections we give a detailed outline of these means.

8.2.7.1. Color-Coded Distance Information and Slope

We have developed three different kinds of visually enhancing interesting events of the patient’s parameter development, i.e., the distance of the parameter value from the intended value, the progress of the parameter value from the initial value, when the treatment started, and the rising and falling of the parameter values. Each of these modes is suited to bring to light very different qualities of the data (i.e., the patient parameters in combination with the applied treatment plans). We use color to visually encode this information because the task of comparing different plans and selecting the best working treatment is of major interest. Color is a strong visual attribute which gives an immediate impression of the different extents of success or failure of applied treatment plans. Thus, it facilitates the comparison and appraisal of plans and conditions. Although we are aware of the problem that colors may have different meanings depending on the cultural background, we tried to find most intuitive color-mappings for the western culture. In addition, we decided for discrete color scales to make qualitative changes in the parameter development stand out more clearly. In the following subsections the three kinds of color-coded information for assessing success or failure of a treatment plan and the necessary data abstractions are outlined.

Distance to Intended Value

The distance of the patient’s parameter value (which is affected by the treatment plan) from the intended value is an important information to immediately identify critical values. We have enhanced the visibility of this information by using a color scheme which maps the relative distance of parameter values from the most deviant value within the range of possible parameter values (the minimum or maximum value of the parameter) to the intended value (e.g., 90-92% oxygen saturation) on a sequential color scale [Harrower & Brewer] from dark magenta to light magenta (see Figure 46). The association of more extreme values with more saturated color entails an intuitive mapping (e.g., very saturated magenta indicates very bad values). Moreover, a scale from very saturated to very pale colors has an implicit ordering and can immediately be interpreted by the eye.

Figure 50 shows one treatment plan that has been applied on three different patients, which are vertically aligned. The charts and treatment plans are colored according to the patients’ parameter values’ distance to the intended value. We selected ranges with a big distance to the intended value with the range slider, which draws attention to critical cases. In doing so, the visualization emphasizes the bad tcpSO2 values of the first patient. In contrast, the minor changes in the tcpSO2 curve of the third patient are close to the intended value range and do not draw off the user’s attention.

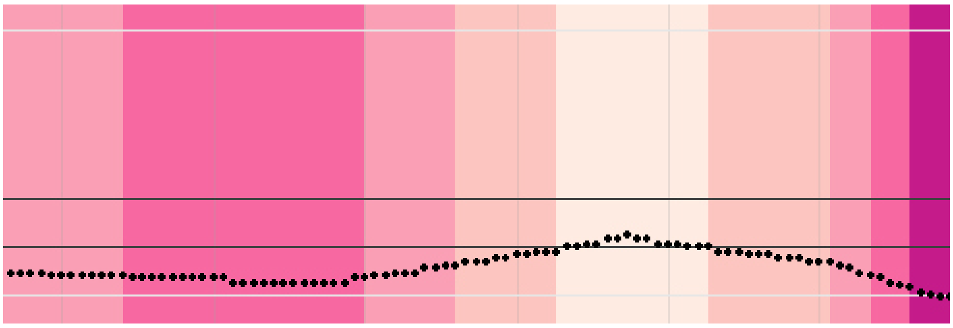

Figure 46: Distance to intended value: Highlighting the distance of the patient's parameter values to the intended value (dark magenta: extreme values, light magenta: inside the intended value range). The range of intended values is indicated by the two dark horizontal lines. This mode helps physicians to identify critical values at the first sight.

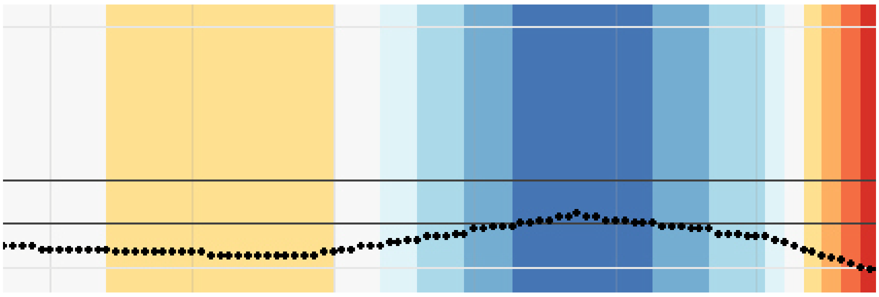

Figure 47: Progress from initial value: Highlighting the progress of the parameter values relative to the initial value when the treatment plan was started (white: start value, dark blue: intended value, dark red: departure from the intended value). This mode shows to what extent the applied treatment plan has the intended effect on the patient's condition.

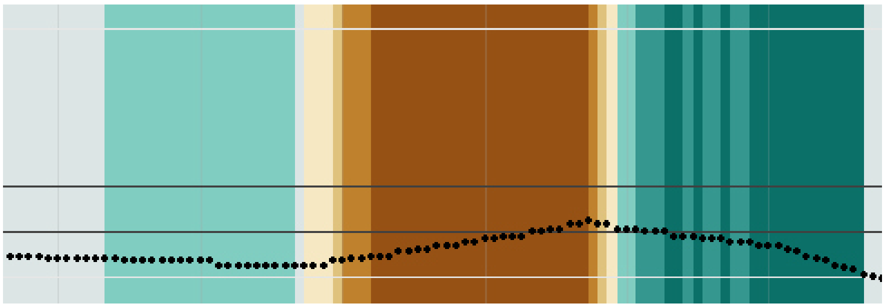

Figure 48: Slope: Highlighting the slope of a parameter value (turquoise: drop, brown: rise). This mode helps to identify the immediate effects of applied clinical actions. For a more robust coloring we take the mean value of seven data points to compute the slope.

Progress from Initial Value

Progress from Initial Value The main criterion for the appraisal of treatment plans is how well the plan succeeds in achieving the intended effects. To visually encode this information, we consider the relative progress from the initial value (e.g., the oxygen saturation at the start of applying the treatment). The initial value is associated with 0% success and the intended value with 100% success of the treatment. In this way, success is expressed by the extent to which the oxygen saturation progresses from the initial value towards the intended value while applying the plan. On the other hand, if the patient's oxygen saturation moves into the opposite direction the success of the treatment plan is negative, which means that the treatment plan fails in achieving the intended effects.

This abstraction method (namely, calculating the progress of the parameter values relative to the initial value) has the advantage of normalizing this measure for different parameters with different scales. Moreover, this measure represents the actual effect of the applied treatment plan, no matter what the initial value was.

It is important to define the success of a treatment with respect to the initial parameter value which is illustrated by the following example: An increase of oxygen saturation from 80% to 85% implies that the applied treatment plan has failed if the intended value was 90%. An increase from 85% to 90% means success for the applied treatment plan, even though in both cases the oxygen saturation was increased by 5%. In the first example, the 5% represent a mere 50% of the distance between initial value (80) and intended value (90). In the second, however, the 5% stand for 100% of the distance between initial value (85) and intended value (90). Absolute values (e.g., oxygen saturation was raised by 5%) fail in communicating success or failure of a plan, if no further information is given.

To visually communicate the different qualities of progress of applied treatment plans we use a diverging color scale [Harrower & Brewer] from blue over white to red. We map the success values from 0% (value at plan start) to 100% (intended value) on colors from white to blue. In case the patient's values move into the opposite direction, the values from 0% to -100% are mapped on a color range from white to red. In this way, blue means good, red means bad, and, again, more saturated colors represent more extreme values (see Figure 47).

Both, the distance to the intended value and the progress from the initial value represent the patient's condition a nd the success of a plan with respect to one specific parameter. So far, this is only an enhancement of information that is also given by the line chart. However, when it comes to getting an immediate impression of a patient's condition that depends on several different parameters, or to assessing the success of treatment plans which may intend to affect two or more parameters (e.g., 'stabilize the patient's tcpSO2 between 90% and 92% and the patient's PCO2 between 40 and 60 mmHg'), things get difficult: data can no longer be retrieved easily from the raw individual parameter curves. The overall success of the plan as well as the overall condition of the patient have to be determined. To this end, we stick to the rules of the infinite-valued Gödel logic [Gödel, 1932] with the following truth degree set W∞:

$$W_\infty = [0,1] = \{x \in \mathbf{R} \mid 0 \leq x \leq 1\}$$

and the following truth degree functions:

$$u \wedge v = min\{u, v\}$$

$$u \vee v = max\{u, v\}$$

$$\neg u = \begin{cases} 1 & \text{if } u = 0 \\ 0 & \text{if } u > 0 \end{cases}$$

$$u \rightarrow v = \begin{cases} 1 & \text{if } u \leq v \\ 0 & \text{if } u > v \end{cases}$$

According to these rules, the overall success of a plan that intends to achieve an effect on parameter A and an effect on parameter B is determined by the minimum of both success values (i.e., success (A and B) = minimum (success (A), success (B)). In turn, the overall success of a plan that intends to achieve a given effect on parameter A or an effect on parameter B is determined by the maximum of both values (i.e., success (A or B) = maximum (success (A), success (B)).

Figure 51 shows the highlighting of the progress of the parameter values from the initial value. The intended effect of the selected treatment plan is to 'stabilize the patient's tcpSO2 between 90% and 92% and the patient's PCO2 between 40 and 60 mmHg', thus, the overall success of treatment plan execution at a given point in time is determined by the minimum of the success value for stabilizing the patient's tcpSO2 and the success value for stabilizing the patient's PCO2. The selected treatment plan is highlighted according to this overall success, while each parameter chart is highlighted according to the success of stabilizing its specific values. As we can see, towards the end of the focus region the execution of the treatment plan has a positive effect on the patient's PCO2, while the patient's tcpSO2 values are still too high. This results in a quite low overall success rate of the plan.

Slope

Visually enhancing the slope of the parameter chart, i.e., if the parameter values rise or drop, reveals some interesting insights as well. In particular, it eases the identification of an immediate correlation of applied clinical actions and changes of the parameter value. In contrast to the other two highlighting modes, we don't encode 'good' and 'bad' changes (in terms of moving towards the intended value) in the parameter chart. This is why we decided for a diverging color scheme [Harrower & Brewer] that is as neutral as possible: we represent negative slopes with turquoise highlighting and positive slopes with brown highlighting (see Figure 48).

8.3. Interactive Step-Wise Exploration of Treatment Plans, Patient Data, and Effects

Besides the highlighting of color-coded information, the step-wise investigation of effects of applied treatment plans and clinical actions is supported by the following interactive means: filtering for curve events of interest, aligning plans or actions below each other to ease comparison, and selecting regions of interest by means of a focus window. In the following subsections we give a detailed outline of these means.

8.3.1. Filtering Color-Coded Information

The user may be interested in only specific curve events, for instance, critical values or extreme drops of the curve. Thus, we provide a range slider to filter for such events of interest which makes them stand out immediately (see Figure 49).

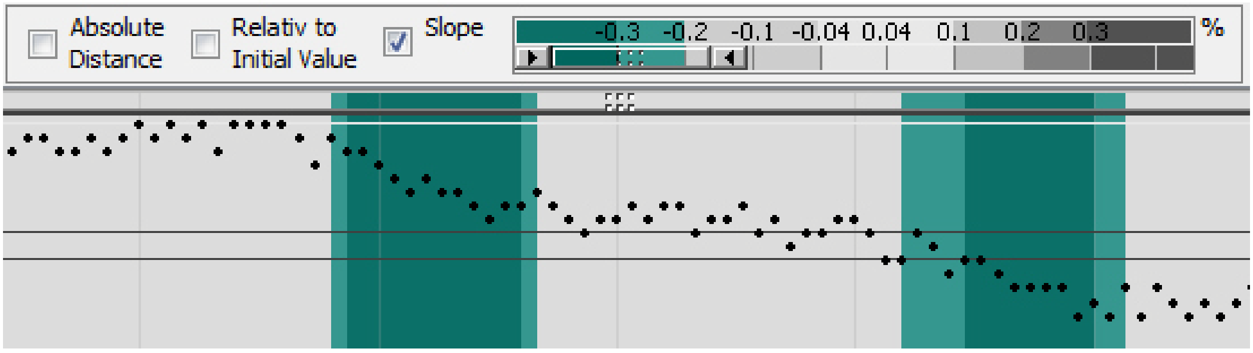

Figure 49: Selecting only extreme drops of the parameter chart with the range slider makes the coloring of minimal drops and rises of the chart disappear.

8.3.2. Aligning Plans or Actions

To facilitate comparison between the effects one treatment plan had on two different patients or the effects of two alternative treatment plans we provide a tool to align selected treatment plans below each other together with their parameter curves (see Figure 50). The color highlighting of these plans gives an immediate impression of differences in how well the plans have worked. Similar to the comparison of treatment plans, we provide means to compare all instances of applying a specific clinical action. This allows the user to investigate how well this specific clinical action works (e.g., the application of a certain drug). The effects of applying this action once may not be representative; also, looking for all applications of this action and subsequently comparing them without visual assistance can be quite a laborious task. Hence, when selecting a clinical action of interest, all instances of applying this action are highlighted; together with the parameter chart they are aligned below each other to facilitate the comparison of the effects the applications of this action have had on the patient's parameters (see Figure 52).

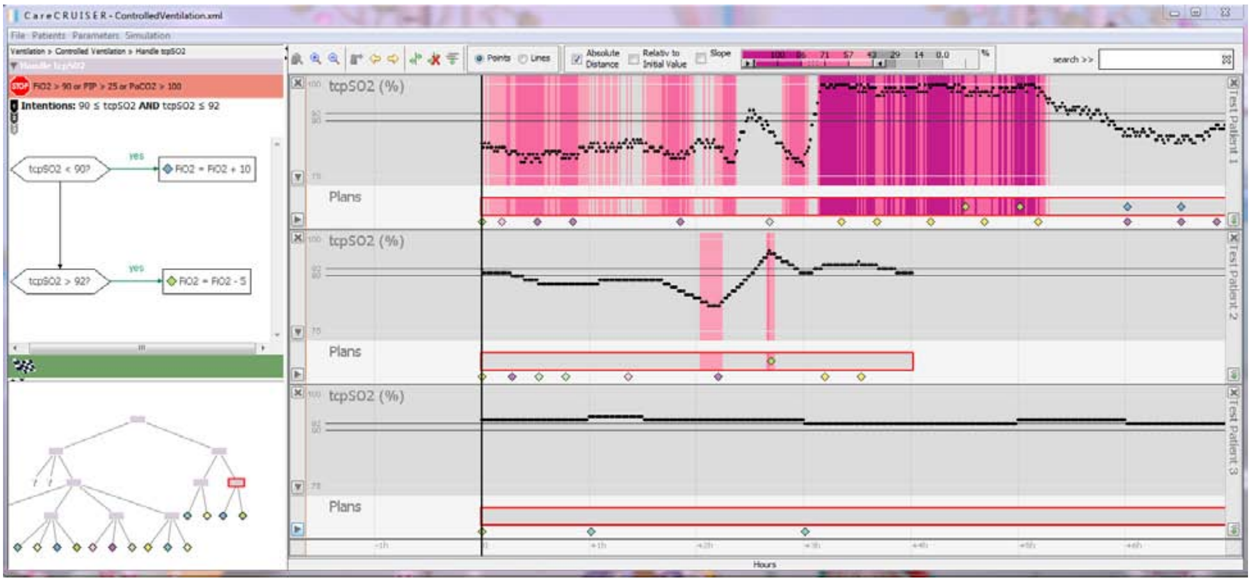

Figure 50: The same treatment plan has been applied on three different patients (aligned vertically). The charts and treatment plans are colored according to the color scheme of the parameter values' distance to the intended value. Selecting ranges with big distance to the intended value with the range slider draws attention to critical cases and brings out the differences between the conditions of the three patients.

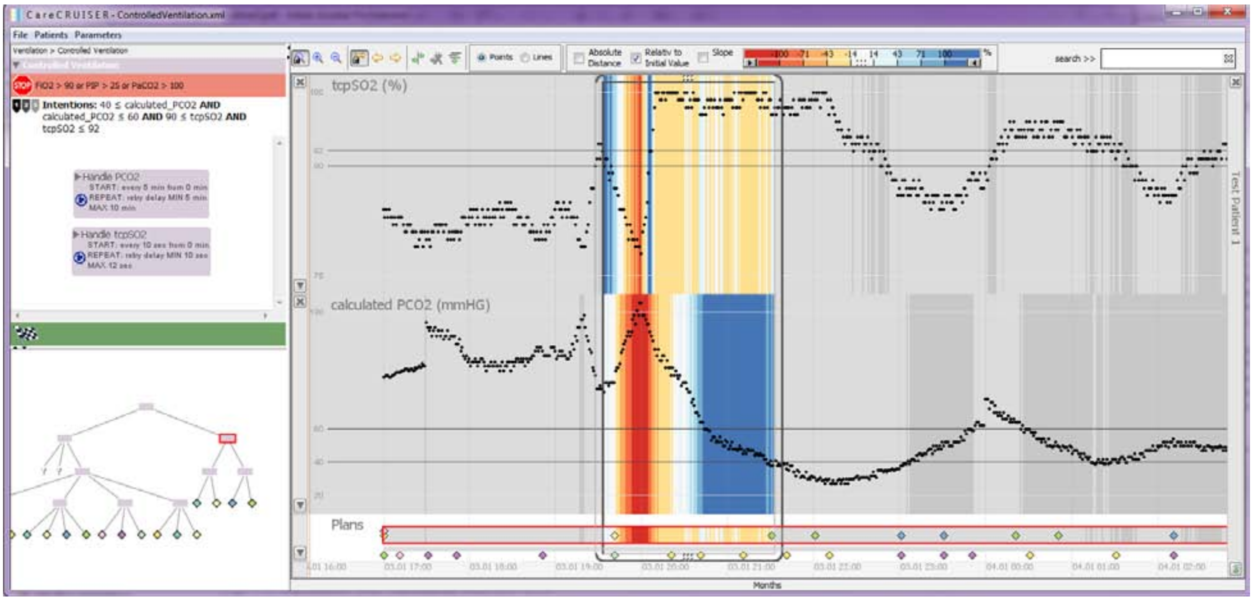
CareCRUISER - ControlledVentilation.xml
File Patients Parameters
Ventilation > Controlled Ventilation
FiO2 > 90 or PIP > 25 or PaCO2 > 100
Intentions: 40 ≤ calculated_PCO2 AND calculated_PCO2 ≤ 60 AND 90 ≤ tcpSO2 AND tcpSO2 ≤ 92
Handle PCO2
START: every 5 min from 0 min
REPEAT: retry delay MIN 5 min MAX 10 min
Handle tcpSO2
START: every 10 sec from 0 min
REPEAT: retry delay MIN 10 sec MAX 12 sec
Points
Lines
Absolute Distance
Relativ to Initial Value
Slope
-100 -71 -43 -14 14 43 71 100
search >>
tcpSO2 (%)
calculated PCO2 (mmHG)
Plans
Test Patient 1
Months

Figure 51: Highlighting the progress of parameter values from the initial value towards the intended value, results in getting an immediate idea of success or failure of the applied treatments. The overall success of executing the selected treatment plan at a given point in time is determined by the minimum of the success value for stabilizing the patient's tcpSO2 and the success value for stabilizing the patient's PCO2. The low success values for stabilizing the patient's tcpSO2 lead to a low overall success for the selected treatment plan.

8.3.3. Focus Window

After color-highlighting and filtering interesting curve figure information and aligning plans or actions of interest, the user may investigate a specific region (e.g., each hour after the different applications of one clinical action). To support this task we have developed a focus window which grays out the color-information outside its borders. The width of the window is varied by dragging its sideborders while the position of the window is varied by dragging its top- or bottom-border horizontally to the desired location. Thus, the user can drag the window over the display area of interest and look for noticeable vertical color-patterns. Extreme curve events outside the focus window are indicated by slightly darker regions which maintains some context information.

Example: Discovering Less Efficient Treatment Choices

When assessing our new methods with our test datasets we could reveal an interesting pattern. The intention of the 'Handle tcpSO2' plan is to stabilize the patient's tcpSO2 between 90% and 92% by adjusting the supply of FiO2. We selected the clinical action 'reduce the FiO2 when the patient's oxygen saturation rises above 92%' and vertically aligned all applications of this action. When highlighting the slope of the patient parameter tcpSO2 and filtering for negative slopes (i.e., the oxygen saturation drops), we can see that almost every time this action was applied the oxygen saturation dropped after some delay. In the meantime the action was applied again which caused another drop of the oxygen saturation below 90%. Therefore, the complementary action had to be applied ('increase the FiO2 when the patient's oxygen saturation drops below 90%'). Helping physicians to reveal this behavior leads to more effective treatment, a reduction of treatment costs, and a better stabilization of the patient (see Figure 52).

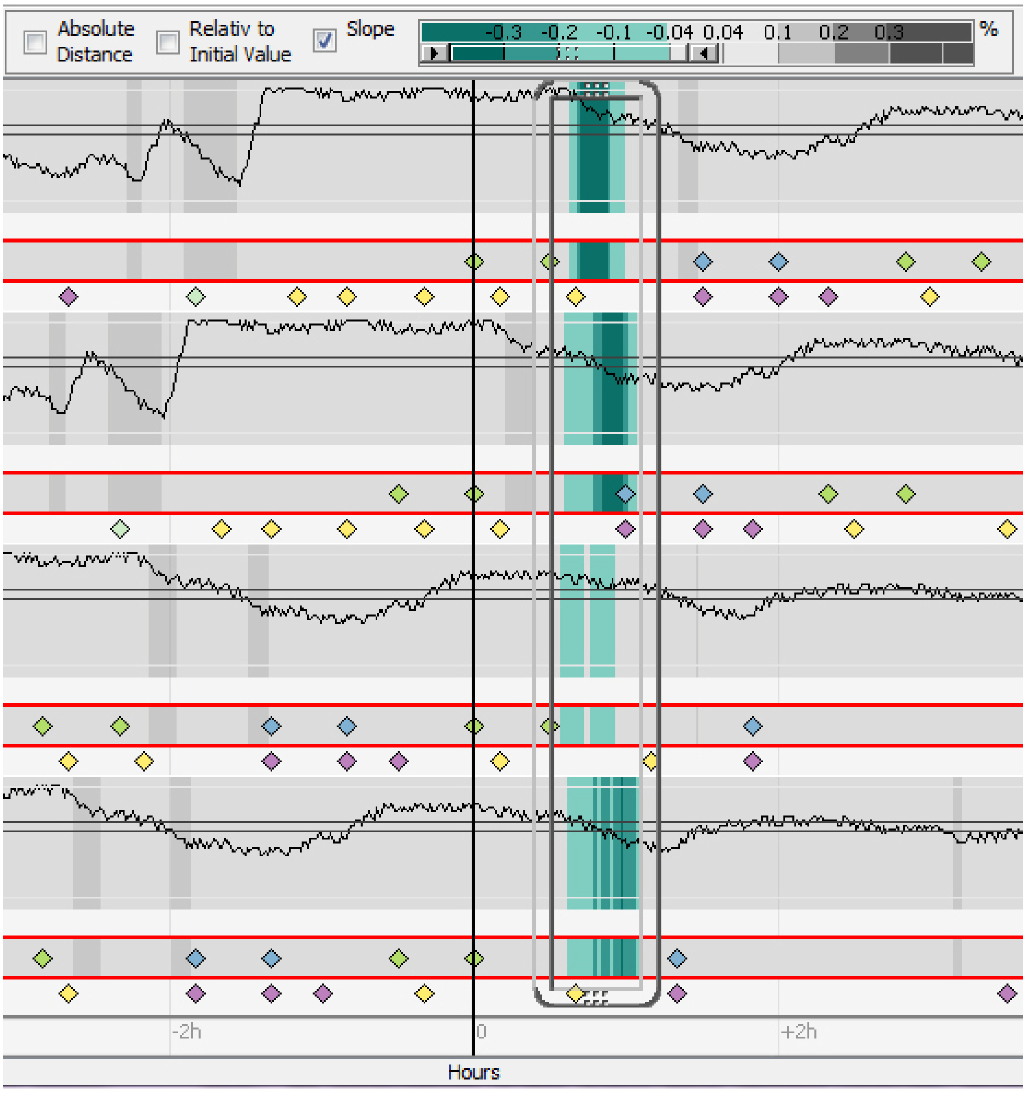

Figure 52: Revealing the delayed drop of the patient's tcpSO2 when applying the clinical action 'reduce the FiO2 when tcpSO2 rises above 92%'. All applications of this action were aligned vertically along the black line; the negative slopes of tcpSO2 were highlighted. Dragging the focus window over the time span after applying the action reveals a vertical turquoise pattern (drops of the tcpSO2 curve) with some delay to the application of the action.

8.3.4. Investigation of Multiple Patients

The visualization allows for the investigation of two or more patients at the same time. The available display height is equally distributed to the patient facets (see Figure 50). However, when the patient facets reach a defined minimum height, the view gets scroll-able. To this end, collapsed patient facets allow for investigating a bigger number of patients at the same time (see **Figure 53**). These facets can be expanded, collapsed or closed individually or collectively. Mouse-over windows allow for the identification of the patient and provide additional demographic information.

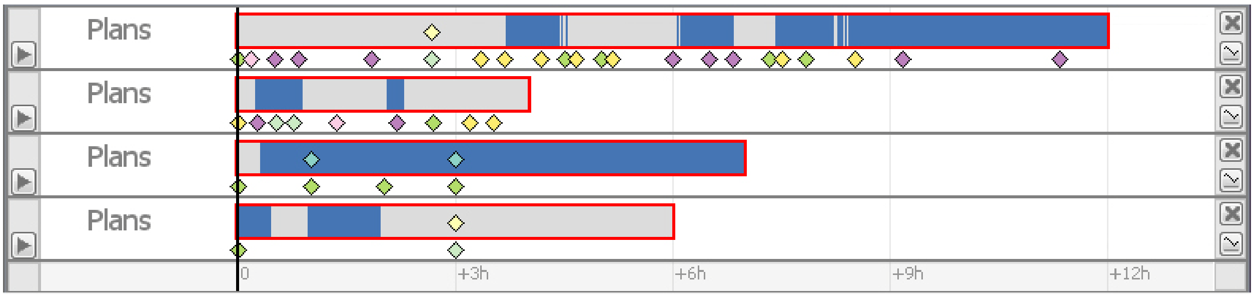

Figure 53: Collapsed patient facets. We have aligned a given treatment plan that has been applied to four different patients. When highlighting the parameter values' progress from the initial value and filtering for parameter values within the intended range (blue) one can still tell that the treatment plan succeeded in stabilizing the values of the first and third patient while failed for the second and fourth patient.

8.3.5. Other Interactive Means

Apart from the interactive means we have outlined in the previous section, we provide minor interactive means which enable smooth workflow, customized perspectives, as well as easy access to diverse kinds of information.

- **Selection of patient parameters.** Besides the parameters specified within the treatment plan's intentions, additional parameters of interest cannot be generalized and highly depend on the specific condition of the patient. Therefore, we provide means to add and remove patient parameters to/from the visualization (see dropdown-box at the right hand pane in Figure 37). Thus the user is able to create customized perspectives of the visualization.

- **Highlighting of clinical actions.** A clinical action of interest can be selected from a complete list of actions that have been applied to the patient (see dropdown-box at the right hand pane in Figure 37). Subsequently, all applications of this action are highlighted within the applied treatment plans. This rids the user of the laborious task of looking for events of this action by moving the mouse to every single action that has been applied in order to read the tool-tip labels.

- **Details on demand.** Since the visual encoding of all information available would lead to a very crowded and cluttered representation, decisions as to which information should be visualized have to be weighed carefully with respect to the major objectives of the visualization. Still, the information that is not chosen for visualization is important and needs to be easily accessible. To achieve this, we provide detailed information on each element (e.g., plan, action, parameter value, etc.) on demand. Mouse-over tool-tip labels are a very effective way to provide detailed information on elements of interest while not interfering with the clearly arranged visual representation.

- **Navigating the timeline.** Since displaying the whole patient history in one picture does not allow for a detailed exploration, we provide means to move along the timeline. In doing so, the user is able to focus on selected sections of interest in the patient's history as well as to investigate the near future as to which treatment plans and clinical actions are to be applied and what the appropriate temporal constraints are. Thus, this kind of visualization is not only an instrument to investigate the patient's history and to monitor the course of the patient, but also serves as a visual assistance and reminder of applying adequate treatment plans and clinical actions in time.

9. Example Visualizations of ReMINE Protocols

After giving a detailed outline of our visual encodings and interactive means, we give example screenshots of our visualization applied to the ReMINE protocols in the following subsection (see Figure 54 to Figure 59). Since by now we do not have all protocol- and patient-data required to show the full potential of the visualization module, these screenshots are understood as preliminary, basic examples.

We have automatically generated example applications of the protocols containing all sub-plans and clinical action specified within the protocol and executed in sequence (unless defined differently). We have generated these data so that part of the protocol application takes place in the future to show the visualization of both, past and future plans.

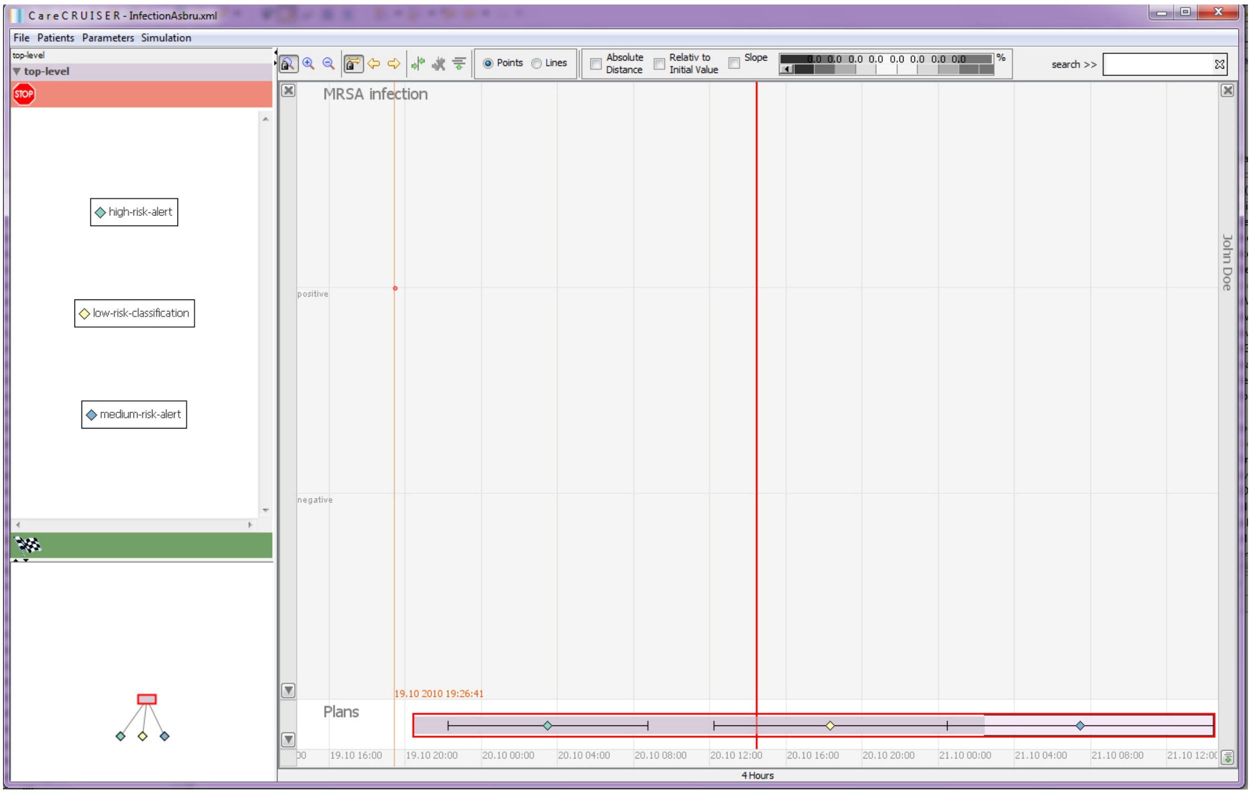

Figure 54: Screenshot of the “InfectionAsbru.xml” protocol. The red line separates the past from the future.

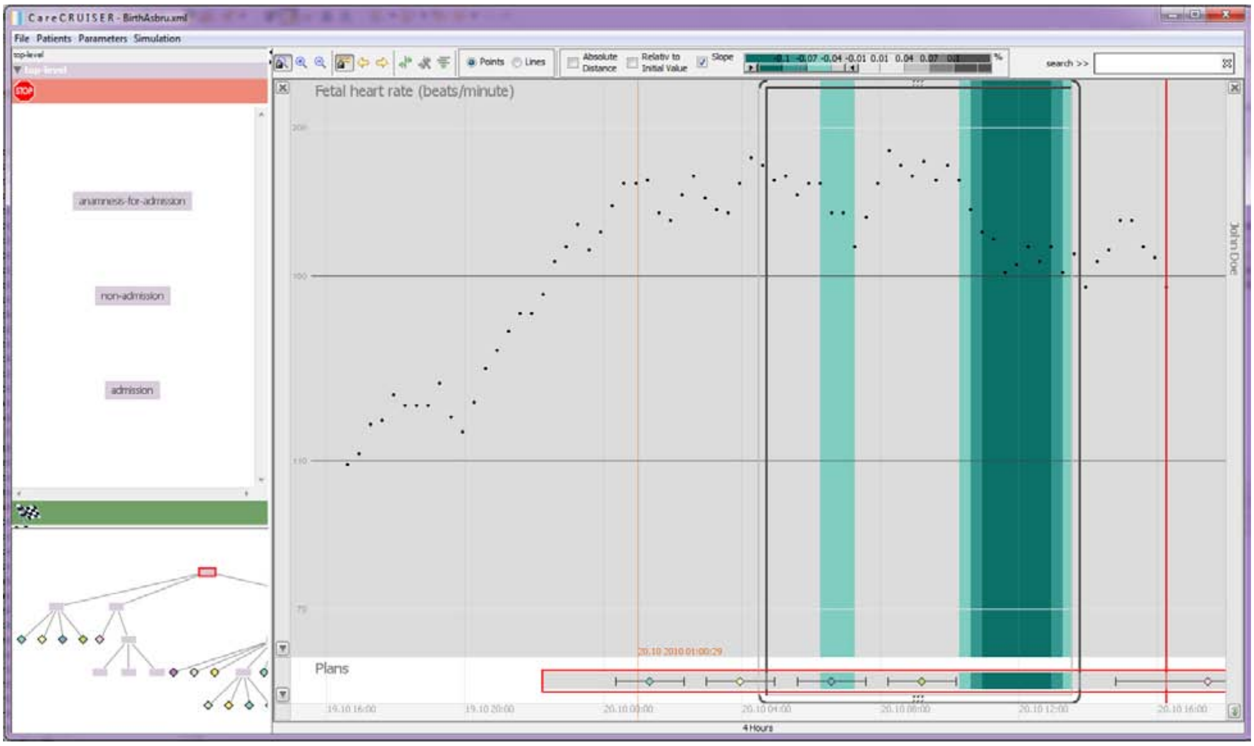

Figure 55: Screenshot of the “BirthAsbru.xml” protocol. Highlighting the negative slope of the parameter chart of the simulated fetal heart rate emphasizes strong downwards trends (for a more robust coloring we take the mean value of seven data points to compute the slope.).

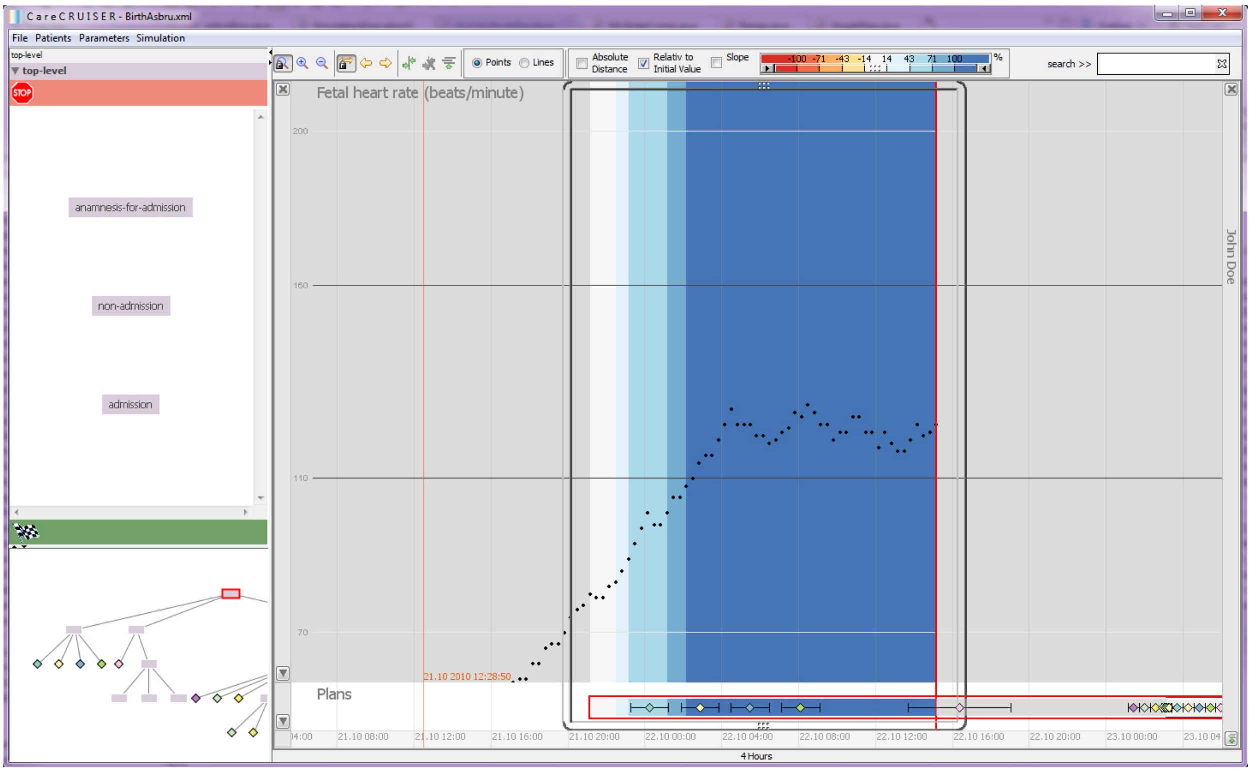

Figure 56: Screenshot of the “BirthAsbru.xml” protocol. Color-highlighting the distance of the parameter values relative to the initial value emphasizes the positive development of the parameter values towards the intended value range during treatment plan execution.

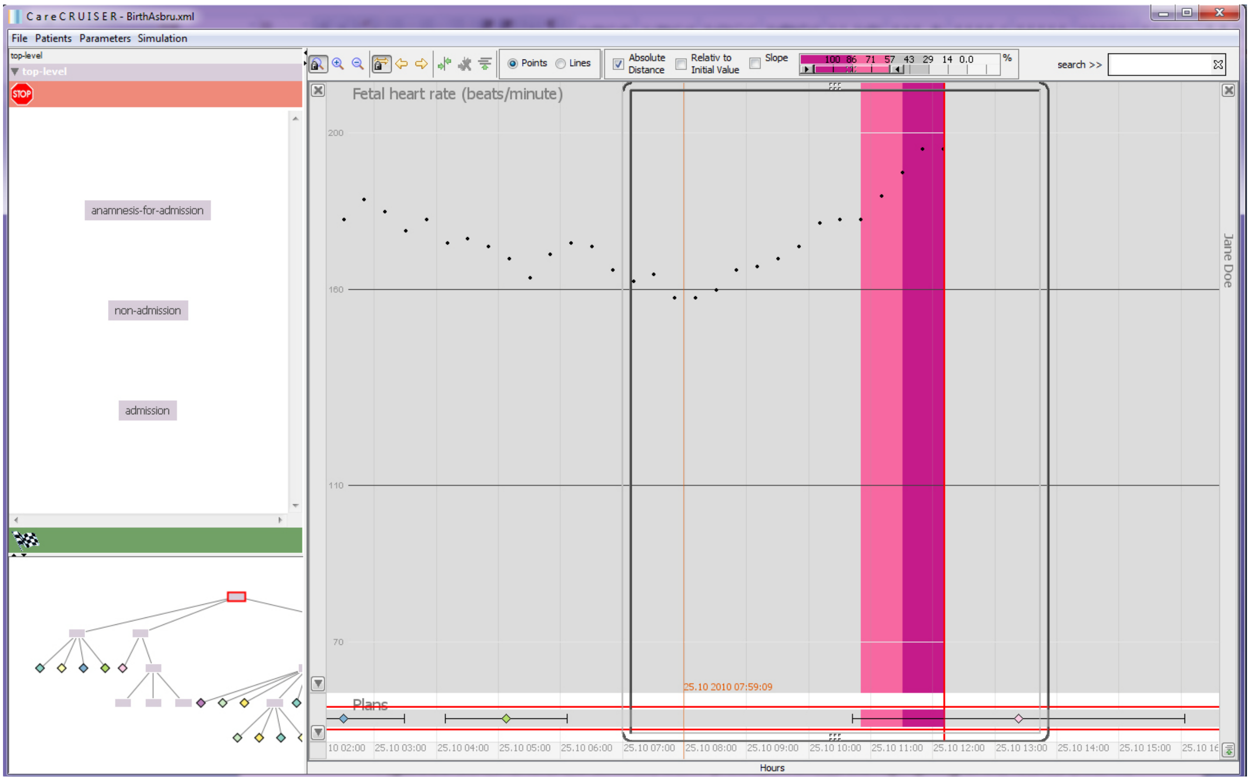

● Points ○ Lines

☑ Absolute Distance □ Relativ to Initial Value □ Slope

Figure 57: Screenshot of the “BirthAsbru.xml” protocol. Color-highlighting the relative distance of the simulated values of the fetal heart rate to the intended value range gives an immediate impression of the fetus’s condition. In this screenshot we use the range slider to filter for values that are at least 57% away from the intended value.

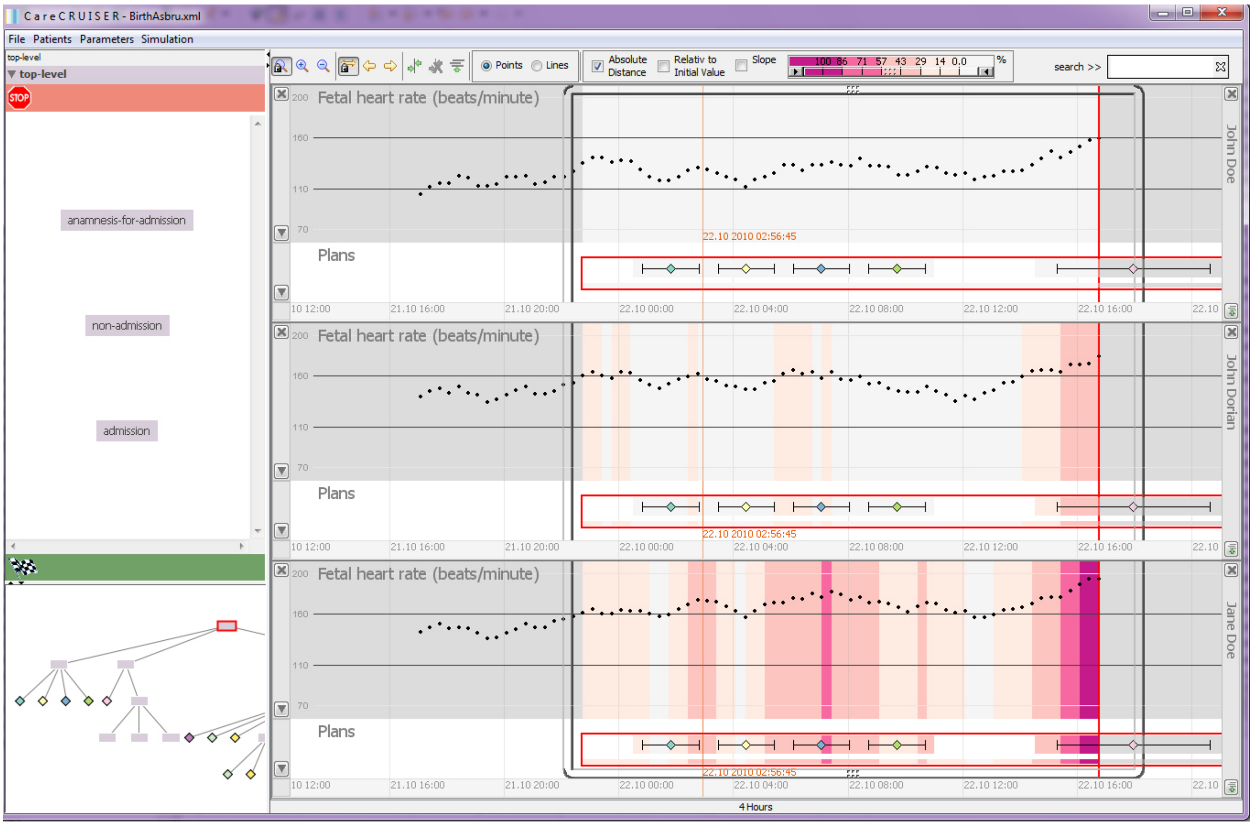

Figure 58: Screenshot of the “BirthAsbru.xml” protocol. It shows the highlighting of the relative distance of the simulated values to the intended value range of three different patients. The patient with the most critical condition can be identified immediately, by the amount of magenta coloring.

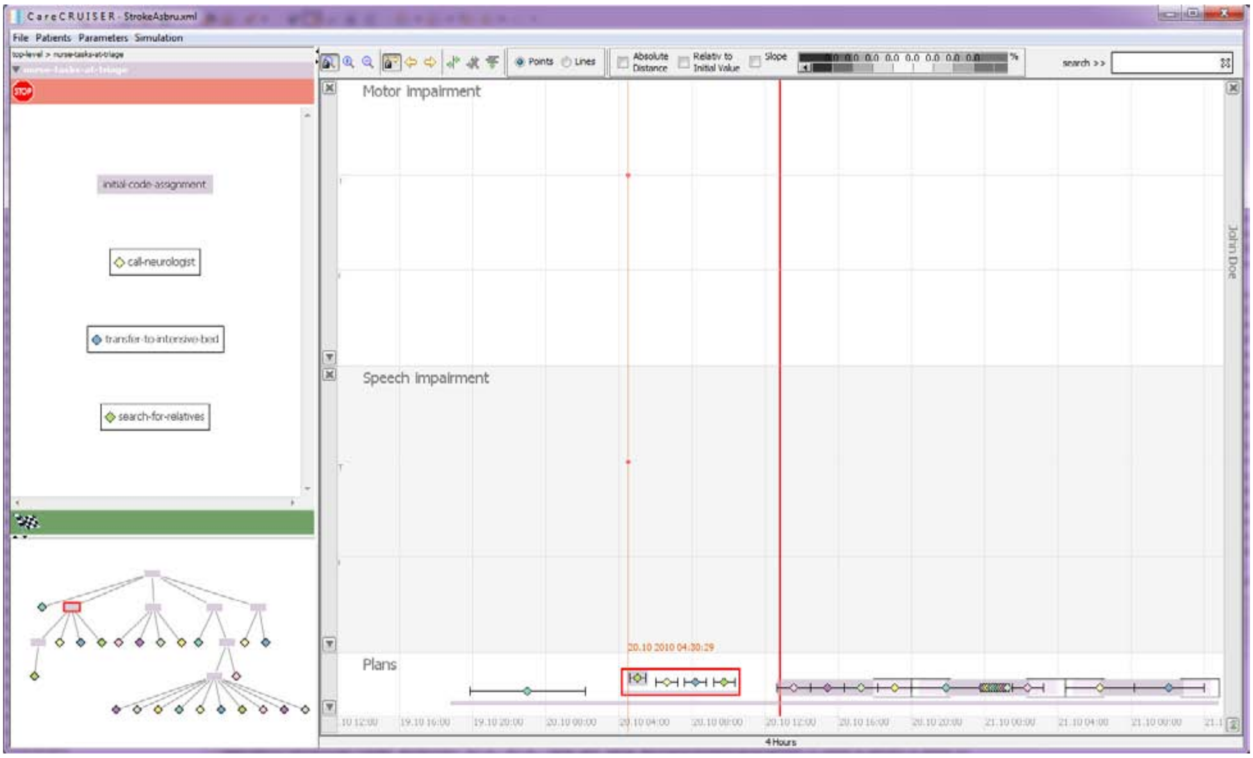

Figure 59: Screenshot of the “StrokeAsbru.xml” protocol.

10. Evaluation

To ensure the quality of our visualization module we conducted a heuristic usability evaluation, on the one hand, and we gathered feedback from a medical expert, on the other hand. In this context we used two different CPGs about ventilating infants together with five data sets of real patient data (four data sets corresponding to the first CPG and one data set corresponding to the second CPG).

To ensure the usability of the prototype we decided for a heuristic usability evaluation according to Forsell and Johansson [Forsell & Johansson]. They presented a best practice set of 10 heuristics out of 63 heuristics (from 6 earlier published heuristic sets):

1. **B5. Information coding.** Perception of information is directly dependent on the mapping of data elements to visual objects. This should be enhanced by using realistic characteristics/techniques or the use of additional symbols.

2. **E7. Minimal actions.** Concerns workload with respect to the number of actions necessary to accomplish a goal or a task.

3. **E11: Flexibility.** Flexibility is reflected in the number of possible ways of achieving a given goal. It refers to the means available to customization in order to take into account working strategies, habits and task requirements.

4. **B7: Orientation and help.** Functions like support to control levels of details, redo/undo of actions, and representing additional information.

5. **B3: Spatial organization.** Concerns users' orientation in the information space, the distribution of elements in the layout, precision and legibility, efficiency in space usage and distortion of visual elements.

6. **E16: Consistency.** Refers to the way design choices are maintained in similar contexts, and are different when applied to different contexts.

7. **C6: Recognition rather than recall.** The user should not have to memorize a lot of information to carry out tasks.

8. **E1: Prompting.** Refers to all means that help to know all alternatives when several actions are possible depending on the contexts

9. **D10: Remove the extraneous.** Concerns whether any extra information can be a distraction and take the eye away from seeing the data or making comparisons.

10. **B9: Data set reduction.** Concerns provided features for reducing a data set, their efficiency and ease of use

This new set is especially tailored to the evaluation of common and important usability problems in Information Visualization techniques.

It is commonly assumed that three to five expert evaluators are sufficient for a heuristic usability evaluation [Holzinger, 2005]. Thus, we conducted the study with four evaluators, who have considerable knowledge about usability principles.

To ensure unbiased evaluations, we performed separate testing sessions with each evaluator. An observer was present at each testing session to answer questions about the domain and to give hints when the evaluator was in trouble (which was not necessary at any time of the four testing sessions). Each evaluator went through the visualization two times. The first round was aimed at getting a feeling for the flow and the general scope of the visualization. In the second round the evaluator was supposed to focus on visual and interactive interface elements with respect to the given list of heuristics [Forsell & Johansson][Forsell & Johansson]. Each evaluator was asked to solve a list of tasks by means of the visualization module:

1. to find out how the patient's condition changes during the execution of the treatment plan;

2. to find out which clinical actions were applied to the patient in the context of treatment plan execution and when these actions were applied;

3. to find out which effects the clinical actions have on the patient's condition;

4. to identify critical parameter values in the course of treatment plan execution) and to note down the found problems with reference to the violated usability principles.

Moreover, the evaluators were asked to rate the severity of the problem (1 representing the lowest severity and 5 the highest).

This evaluation resulted in a list of 32 usability problems. The majority of these problems relate to the two usability principles 'orientation and help' and 'consistency'. 'Orientation and help' deals with means such as support to control levels of details, redo/undo of actions, and representing additional information, while 'consistency' refers to the way design choices are maintained in similar contexts, and are different when applied to different contexts [Forsell & Johansson].

Most of the found problems can easily be fixed, for instance, by adding tooltip text to buttons. On the other hand, the evaluation also revealed more difficult problems, like the lack of a possibility to align the grid with the clinical actions. Moreover, in case a clinical action was applied multiple times and the user selects one instance of this clinical action, all instances of application of this action are highlighted and there is no visual difference between the one instance the user has selected and the other instances of applying this action. Another interesting problem that should be mentioned here is that the evaluators would like to have a history of view modifications and jump back and forth between views while skipping a number of modification steps.

In addition, we assessed the benefits and shortcomings of our visualization in collaboration with a physician. This led to some interesting insights and new ideas. The physician appreciated the possibility to assess the patient's condition at the time of treatment plan application in immediate combination with the comprehensive information about these treatment plans.

Our discussions led to some new ideas about the color schemes. By now a given distance below the intended value is encoded with the same color as the given distance above the intended value. In medical care, however, these two cases may not be equally weighted in terms of critical conditions. Moreover, depending on the given parameter, variability in the upper range of the chart may indicate more harm than in the lower range. Thus, we think about integrating more detailed semantic information about each parameter and adjusting the color scheme to represent this information (e.g., more critical colors for values in the upper range of the chart). In addition, using one color to represent values above the intended range and another color to represent values below this range would add one dimension of information. We will consider these ideas and systematically assess different color schemes to find the most adequate.

Another valuable idea stemming from this collaboration is to separate the means for assessing a patient's condition from those for appraising the quality of the treatment recommendations. Commonly, these tasks are carried out at different points in time and maybe even by different persons. Means for assessing a patient's condition include the coloring and filtering of distances of parameter values from the intended values, and the slope of the parameter chart. Appraising the quality of treatment plans is supported by coloring and filtering the the patient's parameter values' progress from the initial value as well as the slope of the parameter chart in correlation with applied clinical actions. To separate these user views will reduce the quantity of information provided at once.

However, we will carefully consider all found problems in order to improve the visualization.

11. Integration with the ReMINE Risk Manager Interface

The visualisation module supplied by TUW is implemented as an executable .jar file. It receives its inputs in a set of files as described in this section.

11.1. Protocols and Guidelines

Protocols and guidelines need to be modeled in Asbru and provided as .xml files. In addition, a .config file with the same file name as the protocol or guideline has to be provided and located in the same directory. This file must contain the IDs of patients who receive treatment according to this protocol and the path to the directory in which the data of these patients is located. This information is given in this format:

```
patients=testpatient1,testpatient2
patients.path=C:\\data\\patients\\
```

11.2. Patient Data

Patient data is given as a .txt file with the file name “<patientID>_<ddmmyyyy>_<hhmmss>_parameters.txt”, for instance:

```
patient1_27092011_163022_parameters.txt
```

The time stamp in the file name indicates the time of the last parameter value measurement. This file provides continuous patient parameter values. It must contain a column “Date” in the form “ddmmyyyy”, and a column “Time” in the form “hh:mm:ss” giving the time of the parameter value measurement. Moreover it contains two columns for each patient parameter. The first column names the patient parameter in the first line. Further lines contain the measured values for this parameter; the second column consists of the name of the first parameter with the suffix “_unit”. It provides the unit of this patient parameter. For instance:

```
Date  Time     Parameter1  Parameter1_unit   Parameter2  Parameter2_unit
22092010    18:24:11    34    mg/dl 80    %
22092010    18:44:11    37    mg/dl 81    %
22092010    19:04:11    40    mg/dl 82    %
```

All columns are separated by tabulator (TAB character). There must not be an additional blank line after the last row of the table.

In order to visualize an appropriate set of parameters, the visualization module requires a .config file specifying which columns of the parameter text file are to be visualized right from the start (others can be added to the visualization interactively) and which columns of the parameter text file are not to be visualized as parameter charts. The name of the file must be “parameters.config” and must provide the information in this notation:

```
noparam=Date,Time,Parameter1_unit,Parameter2_unit
visibleparam=Parameter1
```

Moreover, demographic information about each patient is provided by a .txt file. The file name must be in the form “<patientID>_demographics.txt”, for instance:

```
patient1_demographics.txt
```

This file must contain the ID of the patient, the name of the patient, the age of the patient, the weight, and the sex of the patient. This information is expressed in the following notation:

```
id=patient1
name=John Doe
age=3
ageunit=years
weight=40
weightunit=kg
sex=m
```

These files (i.e., `patient1_27092011_163022_parameters.txt`, `parameters.config`, and `patient1_demographics.txt`) must be located in the directory that is identified by “`patients.path`” in the .config file of the protocol.

11.3. Applied Clinical Actions

A file “<patientID>`_actions.txt`” contains information about treatment plans and clinical actions that have been applied or will be applied on the patient no matter if this plan / action was intended by the Asbru model or not. This file must be located in the directory that is identified by “`patients.path`” in the .config file of the protocol.

This file contains the real start date of the treatment plan / clinical action (column “`RStartDate`”) in the form “ddmmyyyy”, the real start time of the treatment plan / clinical action (column “`RStartTime`”) in the form “hh:mm:ss”. Accordingly, the file contains the real end date of the treatment plan / clinical action (column “`REndDate`”) in the form “ddmmyyyy” and the real end time of the treatment plan / clinical action (column “`REndTime`”) in the form “ddmmyyyy”. The earliest start date (column “`EStrartDate`”), the earliest start time (column “`EStartTime`”), the earliest end date (column “`EEndDate`”), the earliest end time (column “`EEndTime`”), the latest start date (column “`LStartDate`”), the latest start time (column “`LStartTime`”), the latest end date (column “`LEndDate`”), and the latest end time (column “`LEndTime`”) are given in the same form. Start and end time can be the same if the action takes place at a single point in time. In case the real start and end of a treatment plan / clinical action are specified, the other time annotations are either given according to the actual annotations or are substituted by the real start and end time.

Moreover, the ID of the treatment plan / clinical action is given in column “`ID`”. The ID relates to the name of the plan as modeled in the Asbru file (e.g., `<plan name="top-level">`). In case a clinical action was performed that was not intended by the Asbru protocol, the ID is “`-`”.
A column “`Regular`” specifies if the treatment plan / clinical action was intended by the Asbru protocol; in this case it contains the character “`T`” (otherwise “`F`”).

The column “`Action`” contains the character “`T`” if it is a clinical action (i.e., a user performed treatment plan) and “`F`” if is a treatment plan (not user performed).

The column “`Description`” contains an optional description of the treatment plan / clinical action. If no description is provided it gives either the character “`-`” or the dummy description “`description`”.

The following columns describe applied clinical actions in more detail. In case the table row corresponds to a treatment plan (not user performed) these columns contain the character “`-`”.

The column "`Substance`" contains the name of the substance that was given to the patient (e.g., the name of a drug or other kinds of substances, for instance, oxygen). If no substance was given it contains the character "`-`".

The column "`Amount`" gives the total amount of the substance that was given to the patient. In case no substance was given this column contains the character "`-`".

The column "`Unit`" gives the unit of the administered substance ("`-`" otherwise).

The column "`Action_kind`" describes the kind of action, for instance, "`initial setting`", "`increase`", "`decrease`", or "`-`".

In case the substance is administered on a regular basis, the column "`Change`" provides the difference of the total amount of the administered substance to the amount administered before (e.g., "`+2`", "`-10`").

These data columns help to provide more detailed information about all applied clinical actions weather intended by the Asbru protocol or not. Thus, they are used by the visualization to provide comprehensive information about the clinical situation.

11.4. Additional Information about Parameters

In addition, the visualization module requires information on the risk-levels of patient parameters that are subject of the treatment. File "`risklevels.config`" contains this information in the following notation:

```
Parameter.names=Parameter1,Parameter2
#Parameter1
Parameter1.min=0
Parameter1.max=100
Parameter1.low=70
Parameter1.high=80
#Parameter2
Parameter2.min=75
Parameter2.max=100
Parameter2.low=90
Parameter2.high=92
```

At which <parametername>`.min` is the minimal possible value, <parametername>`.max` is the maximal possible value, <parametername>`.low` is the lower limit, and <parametername>`.high` is the upper limit for the normal value of the parameter. This file must be located in the same directory as the protocol.

11.5. Example Data

In this section we provide an example of the required files. It deals with the "`MedicalAsbru.xml`" protocol which is located in the directory "`F:/ReMINE/protocols/`".

The file "`MedicalAsbru.config`" (located in the directory "`F:/ReMINE/protocols/`"):

```
patients=simulation_patient
patients.path=F:\\ReMINE\\patients\\
```

The file “`simulation_patient_28092010_130246_parameters.txt`” (located in the directory “`F:/ReMINE/patients/`”):

```
Date  Time  Example_Parameter Example_Parameter_unit
27092010    13:22:46    34    mg/dl
27092010    13:42:46    37    mg/dl
27092010    14:02:46    40    mg/dl
27092010    14:22:46    41    mg/dl
27092010    14:42:46    43    mg/dl
27092010    15:02:46    40    mg/dl
27092010    15:22:46    40    mg/dl
27092010    15:42:46    40    mg/dl
27092010    16:02:46    39    mg/dl
27092010    16:22:46    35    mg/dl
…
…
…
28092010    13:02:46    74    mg/dl
```

The file “`parameters.config`” (located in the directory “`F:/ReMINE/patients/`”):

```
noparam=Date,Time,Example_Parameter_unit
visibleparam=Example_Parameter
```

The file “`simulation_patient_demographics.txt`” (located in the directory “`F:/ReMINE/patients/`”):

```
id=simulation_patient
name=John Doe
age=3
ageunit=years
weight=40
weightunit=kg
sex=m
```

The file “`simulation_patient_actions.txt`” (located in the directory “`F:/ReMINE/patients/`”):

```
RStartDate RStartTime REndDate    REndTime    EStartDate EStartTime
      EEndDate    EEndTime    LStartDate LStartTime LEndDate
      LEndTime    ID    Regular     Action      Description Substance
      Amount      Unit  Action_kind Change
27092010    19:02:46    29092010    07:02:46    27092010    13:02:46
      29092010    01:02:46    28092010    01:02:46    29092010
      13:02:46    top-level   T     F     description -     -     -     -
      -
27092010    19:47:46    28092010    00:17:46    27092010    19:02:46
      27092010    23:32:46    27092010    20:32:46    28092010
      01:02:46    arrival     T     F     description -     -     -     -
      -
```

```
28092010    01:47:46    28092010    06:17:46    28092010    01:02:46
      28092010    05:32:46    28092010    02:32:46    28092010
      07:02:46    inspection-and-drug-order    T     F     description-
      -     -     -     -
28092010    07:47:46    28092010    12:17:46    28092010    07:02:46
      28092010    11:32:46    28092010    08:32:46    28092010
      13:02:46    nurse-records-prescription  T     T     description-
      -     -     -     -
…
…
…
29092010    07:47:46    29092010    12:17:46    29092010    07:02:46
      29092010    11:32:46    29092010    08:32:46    29092010
      13:02:46    monitoring-of-continuity-after-release  T     T
      description-     -     -     -     -
```

The file “`risklevels.config`” (located in the directory “`F:/ReMINE/protocols/`”):

```
Parameter.names= Example_Parameter
#Example_Parameter
Example_Parameter.min=0
Example_Parameter.max=100
Example_Parameter.low=70
Example_Parameter.high=80
```

For the corresponding visualization see **Figure 60.**

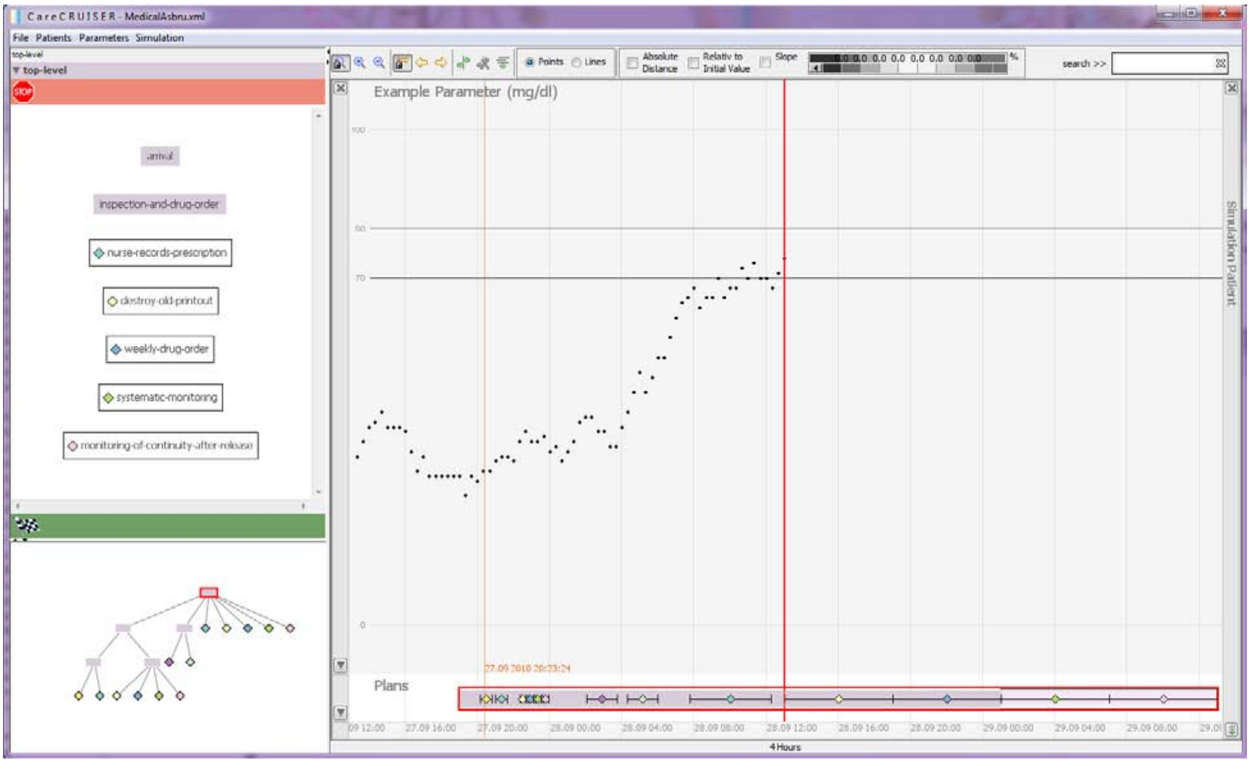

Figure 60: Screenshot of the visualization corresponding to the data files above.

12. Conclusion

This deliverable presents an interactive visualization prototype customised to the needs of the ReMINE project. First, we have given a detailed analysis of the characteristics of all different kinds of information that have to be visualized. Second, we have given a detailed analysis of existing approaches and arrived at the conclusion that these approaches have different strengths and shortcomings but none of them meets all the requirements which result from the analysis of the data. Third, we have developed a visualization module and discussed all decisions for visual information encoding in detail.

On the one hand, the visualization provides charts of the patient's parameters which are mandatory to assess the patient's condition. On the other hand, it provides sophisticated representations of applied treatment plans as well as of plans that are to be applied in the near future, integrating a variety of diverse kinds of information, such as:

- Logical structure of treatment plans,
- Hierarchical structure of treatment plans,
- Execution order of plans,
- Clinical actions applied in the context of executing the plan,
- Complex temporal constraints of the plan, and
- Effects of applied treatment plans.

Our visualization module takes the process information (i.e., the specified intentions of treatment plans) into account to visualize the processes' effects instead of pure data visualizations. We provide various interaction methods to explore the data, processes, and effects interactively. The visualization offers different views about treatment plan execution and the effects on the patient data and it allows for evaluating multiple patients in parallel. While there are a great number of sophisticated tools available that focus on individual aspects of the application of computer-executable treatment plans, our prototype takes the available methods a decisive step further by providing a working environment that offers an intuitive comprehension and grasp of aspects being of interest in the context of evaluating and executing treatment plans.

Moreover, we provide different means and views which enable the user to interactively explore the data. Thus, the visualization supports the revelation of new insights and the generation of hypotheses. These means are designed with special respect to support the assessment and comparison of treatment plans and consequently to support the identification of the best working therapies. There are:

- Means to highlight and filter specific events of parameter developments,
- Means to align treatment plans or clinical actions together with the corresponding patient parameter values to ease comparison,

- Means to move the focus along the time line to support the detection of interesting event patterns while maintaining context information.
- Means to modify and navigate these views.
- Means to visually compare both details and the results of treatment applied to a great number of patients.

13. References

[Aigner, 2001] Aigner, W. (2001). *Guideline Overview Tool (GOT), Technical Report Asgaard-TR-2001-4.* Vienna, Austria: Vienna University of Technology, Institute of Software Technology and Interactive Systems.

[Aigner et al., 2008] Aigner, W., Kaiser, K., & Miksch, S. (2008). Visualization Methods to Support Guideline-Based Care Management. In A. tenTeije, P. Lucas, & S. Miksch (Eds.), *Computer-Based Medical Guidelines and Protocols: A Primer and Current Trends*, pages 140–159. IOS Press, Health Technology and Informatics.

[Aigner & Miksch, 2006] Aigner, W., & Miksch, S. (2006). CareVis: Integrated visualization of computerized protocols and temporal patient data. Artificial Intelligence in Medicine (AIIM), 37(3):203–218.

[Bade, 2002] Bade, R. (2002). Methoden zur Visualisierung von und Interaktion in zeitbasierten Patientendaten und Behandlungsplänen. Master's thesis, Department of Simulation and Graphics, University of Magdeburg.

[Bade et al., 2004] Bade, R., Schlechtweg, S., & Miksch, S. (2004). Connecting Time-oriented Data and Information to a Coherent Interactive Visualization. *Proc. of the 2004 Conf. on Human Factors in Computing Systems (CHI04)*, pages 105–112. ACM Press.

[Burch et al., 2008] Burch, M., Beck, F., & Diehl, S. (2008). Timeline Trees: Visualizing Sequences of Transactions in Information Hierarchies. Proc. of the Working Conference on Advanced Visual Interfaces (AVI'08). New York, NY: ACM, pages 75-82.

[Card et al., 1999] Card, S., Mackinlay, J., Shneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, San Francisco.

[Chittaro et al., 2003] Chittaro, L., Combi, C., & Trapasso, G. (2003). Data mining on temporal data: A visual approach and its clinical application to hemodialysis. *Journal of Visual Languages and Computing ,* 14(6):591–620.

[Ciccarese et al., 2003] Ciccarese, P., Caffi, E., Boiocchi, L., Halevy, A., Quaglini, S., Kumar, A., et al. (2003). The NewGuide Project: Guidelines, Information Sharing and Learning from Exceptions. *Proc. of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003)*, pages 163–167. Protaras, Cyprus: Springer.

[Ciccarese et al., 2004] Ciccarese, P., Caffi, E., Boiocchi, L., Quaglini, S., Stefanelli, A. (2004). A Guideline Management System. *Medinfo 2004,* 11(Pt 1):28–32.

[Ciccarese et al., 2002] Ciccarese, P., Kumar, A., & Quaglini, S. (2002). New-Guide: a new approach to representing clinical practice guidelines. *Proc. of the Workshop on Advances in Clinical Knowledge Management 5, Cancer Research UK*, pages 15–18. London, UK.

[Fekete et al., 2008] Fekete, J., Wijk, J. J., Stasko, J. T., and North, C. (2008). The Value of Information Visualization. *Information Visualization: Human-Centered Issues and Perspectives, Lecture Notes In Computer Science*, vol. 4950, pages1-18. Berlin-Heidelberg, Germany: Springer.

[Few, 2004] Few, S. (2004). Show Me the Numbers: Designing Tables and Graphs to Enlighten. Analytics Press.

[Forsell & Johansson] Forsell, C. & Johansson, J. (2010). An heuristic set for evaluation in information visualization. Proc. of the International Conference on Advanced Visual Interfaces (AVI '10), pages 199–206, New York, NY, USA: ACM.

[Fox et al., 1998] Fox, J., Johns, N., & Rahmanzadeh, A. (1998). Disseminating medical knowledge: The PROforma approach. *Artificial Intelligence in Medicine*, 14:157–181.

[Gennari et al., 2002] Gennari, J., Musen, M., Fergerson, R., Grosso, W., Crubézy, M., Eriksson, H., et al. (2002). The Evolution of Protege: An Environment for Knowledge-Based Systems Development. *International Journal of Human-Computer Studies,* 58:89–123.

[Gödel, 1932] Gödel, K. (1932). Zum intuitionistischen Aussagenkalkül. *Anzeiger der Akademie der Wissenschaften in Wien, Mathematisch- Naturwissenschaftliche Klasse,* 69:65–66.

[Hadorn, 1995] Hadorn, D. (1995). Use of Algorithms in Clinical Practice Guideline Development: Methodology Perspectives. *Clinical Practice Guideline Development: Methodology Perspectives,* 9(95):93–104.

[Harrower & Brewer] Harrower, M. & Brewer, C. (2003) Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37.

[Hinum et al., 2005] Hinum, K., Miksch, S., Aigner, W., Ohmann, S., Popow, C., Pohl, M., & Rester, M. (2005). Gravi++: Interactive Information Visualization of Highly Structured Temporal Data. *Journal of Universal Computer Science, Special Issue on Visual Data Mining,* 11:1792–1805.

[Holzinger, 2005] Holzinger, A. (2005). Usability engineering methods for software developers. *Communications of the ACM*, 48(1):71–74.

[Horn et al., 1998] Horn, W., Popow, C., & Unterasinger, L. (1998). Metaphor Graphics to Visualize ICU Data over Time. *Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-98), Workshop Notes of the ECAI-98 Workshop*, pages 76–81.

[Klimov & Shahar, 2005] Klimov, D., & Shahar, Y. (2005). A Framework for Intelligent Visualization of Multiple Time-Oriented Medical Records. *Proc. of the American Medical Informatics Association Fall Symposium (AMIA 2005)*, pages 405–409.

[Keim, 2002] Keim, D. (2002). Information visualization and visual data mining. *Transactions on Visualization and Computer Graphics*, 8:1–8.

[Kosara & Miksch, 2001] Kosara, R., & Miksch, S. (2001). Metaphors of Movement: A Visualization and User Interface for Time-Oriented, Skeletal Plans. *Artificial Intelligence in Medicine, Special Issue on Information Visualization in Medicine*, 22 (2):111–131.

[Lamping et al., 1995] Lamping, J., Rao, R., & Pirolli, P. (1995). A Focus Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. Proceedings of the Conference on Human Factors in Computing Systems (CHI'95). Denver, Colorado: ACM, pages 401-408.

[Leong et al., 2007] Leong, T. Y., Kaiser, K., & Miksch, S. (2007). Free and Open Source Enabling Technologies for Patient-Centric, Guideline-Based Clinical Decision Support: A Survey. (A. Geissbuhler, R. Haux, & C. Kulikowski, Eds.) *IMIA Yearbook of Medical Informatics 2007, Methods Inf Med,* 46(Suppl. 1):74–86.

[Microsoft Corporation] Microsoft Corporation. *Mircosoft Health Common User Interface*. http://www.mscui.net/Default.aspx, last accessed: June 09, 2009.

[Miksch et al., 1997a] Miksch, S., Shahar, Y., Horn, W., Popow, C., Paky, F., & Johnson, P. (1997). Time-Oriented Skeletal Plans: Support to Design and Execution. *Fourth European Conference on Planning (ECP'97)*, pages 299–311. Springer Berlin / Heidelberg.

[Miksch et al., 1997b] Miksch, S., Shahar, Y., & Johnson, P. (1997). Asbru: A Task-Specific, Intention-Based, and Time-Oriented Language for Representing Skeletal Plans. In E. Motta, F. Harmelen, & C. Pierret-Golbreich (Eds.), *7th Workshop on Knowledge Engineering: Methods + Languages (KEML-97).* Milton Keynes, UK.

[NHS] National Health Services. England. http://www.nhs.uk/Pages/HomePage.aspx, last accessed: June 09, 2009.

[Orsini, 2006] Orsini, R. (2006). Treemap on Rails - O'Reilly Ruby. Retrieved at: December 11, 2009. http://www.oreillynet.com/ruby/blog/2006/07/treemap_on_rails.html.

[Plaisant et al., 2008] Plaisant, C., Lam, S., Shneiderman, B., Smith, M., Roseman, D., Marchand, G., et al. (2008). *Searching electronic health records for temporal patterns in patient histories: A case study with Microsoft Amalga, Technical Report HCIL-2008-13.* College Park, MD: University of Maryland.

[Plaisant et al., 1998] Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., & Shneiderman, B. (1998). LifeLines: Using visualization to enhance navigation and analysis of patient records. *Proc. of the American Medical Informatics Association Fall Symposium (AMIA 1998)*, pages 76–80.

[Polvani et al., 2000] Polvani, K.-A., Agrawal, A., Karras, B., Deshpande, A., & Shiffman, R. (2000). *GEM Cutter Manual, Technical report.* New Haven, CT: Yale Center for Medical Informatics.

[Powsner & Tufte, 1994] Powsner, S., & Tufte, E. (1994). Graphical Summary of Patient Status. *The Lancet,* 344:386–389.

[Quaglini et al., 2001] Quaglini, S., Stefanelli, M., Lanzola, G., Caporusso, L., & Panzarasa, S. (2001). Flexible guideline-based patient careflow systems. *Artificial Intelligence in Medicine,* 22(1):65–80.

[Randelshofer, 2009] Randelshofer, W. (2009). Visualization of large tree structures. Retrieved at: December 11, 2009. http://www.randelshofer.ch/treeviz/index.html.

[ReMINE] REMINE Consortium. REMINE: Supporting Hospitals in Risk Management. Available from: http://www.remine-project.eu/, last accessed: June 09, 2009.

[Seyfang et al., 2002] Seyfang, A., Kosara, R., & Miksch, S. (2002). *Asbru 7.3 Reference Manual.* Technical Report Asgaard-TR-2002-1, Vienna University of Technology, Institut of Software Technology & Interactive Systems, Vienna, Austria.

[Shahar & Cheng, 1998] Shahar, Y., & Cheng, C. (1998). Knowledge-Based Visualization of Time-Oriented Clinical Data. *Proc. of the American Medical Informatics Association Symposium (AMIA'98)*, pages 155–159.

[Shahar et al., 2006] Shahar, Y., Goren-Bar, D., Galperin, M., Boaz, D., & Tahan, G. (2006). Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artificial Intelligence in Medicine,* 38(2):115–135.

[Shahar et al., 1998] Shahar, Y., Miksch, S., & Johnson, P. (1998). The Asgaard Project: A Task-Specific Framework for the Application and Critiquing of Time-Oriented Clinical Guidelines. *Artificial Intelligence in Medicine (AIIM),* 14:29–51.

[Shahar et al., 2003] Shahar, Y., Young, O., Shalom, E., Mayaffit, A., Moskovitch, R., Hessing, A., et al. (2003). DEGEL: A Hybrid, Multiple-Ontology Framework for Specification and Retrieval of Clinical. In M. Dojat, E. Keravnou, & P. Barahona (Eds.), *Proc. of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003),* pages 122–131. Protaras, Cyprus: Springer Verlag.

[Shiffman et al., 2000] Shiffman, R., Karras, B., Agrawal, A., Chen, R., Marenco, L., & Nath, S. (2000). GEM: A Proposal for a More Comprehensive Guideline Document Model Using XML. *Journal of the American Medical Informatics Association (JAMIA),* 7:(5):488–498.

[Shneiderman, 1992] Shneiderman, B. (1992). Tree Visualization with Tree-Maps: 2-D Space-Filling Approach. ACM Trans. Graph., 11, pages 92–99.

[Society for Medical Decision Making, 1992] Society for Medical Decision Making (1992). Proposal for Clinical Algorithm Standards. *Medical Decision Making,* 12(2):149–154.

[Steele & Fox, 2002] Steele, R., & Fox, J. (2002). *Tallis PROforma Primer – Introduction to PROforma Language and Software with Worked Examples.* Advanced Computation Laboratory, Cancer Research, London, UK.

[Terenziani et al., 2004] Terenziani, P., Montani, S., Bottrighi, A., Torchio, M., Molino, G., & Correndo, G. (2004). The GLARE Approach to Clinical Guidelines: Main Features. *Stud Health Technol Inform,* 101:162–166.

[Tufte, 1983] Tufte, E. (1983). *The visual display of quantitative information.* Cheshire, CT, USA: Graphics Press.

[Votruba et al., 2004] Votruba, P., Miksch, S., & Kosara, R. (2004). Facilitating Knowledge Maintenance of Clinical Guidelines and Protocols. In M. Fieschi, E. Coiera, & Y.-C. J. Li (Eds.), *Proc. from the Medinfo 2004 World Congress on Medical Informatics*, pages 57–61. AMIA, IOS Press.

[Wang et al., 2008] Wang, T. D., Plaisant, C., Quinn, A., Stanchak, R., Shneiderman, B., & Murphy, S. (2008). Aligning Temporal Data by Sentinel Events: Discovering Patterns in Electronic Health Records. *Proc. of the ACM Conference on Human Factors in Computing Systems (CHI 2008)*, pages 457–466.