

Interactive Visualization of Effects of Medical Treatment on a Patient's Condition

DISSERTATION

zur Erlangung des akademischen Grades

Doktorin der technischen Wissenschaften

eingereicht von

Theresia Gschwandtner

Matrikelnummer 9949197

an der
Fakultät für Informatik der Technischen Universität Wien

Betreuung: a.o.Univ.Prof.in Dr.in Silvia Miksch

Diese Dissertation haben begutachtet:

(a.o.Univ.Prof.in Dr.in Silvia
Miksch)

(Univ.Prof. Dr.-Ing. Bernhard
Preim)

Wien, 13.09.2012

(Theresia Gschwandtner)

Interactive Visualization of Effects of Medical Treatment on a Patient's Condition

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktorin der technischen Wissenschaften

by

Theresia Gschwandtner

Registration Number 9949197

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: a.o.Univ.Prof.in Dr.in Silvia Miksch

The dissertation has been reviewed by:

(a.o.Univ.Prof.in Dr.in Silvia
Miksch)

(Univ.Prof. Dr.-Ing. Bernhard
Preim)

Wien, 13.09.2012

(Theresia Gschwandtner)

Erklärung zur Verfassung der Arbeit

Theresia Gschwandtner
Taborstr. 18, 1020 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit - einschließlich Tabellen, Karten und Abbildungen -, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

(Ort, Datum)

(Unterschrift Verfasserin)

Acknowledgements

First of all I want to thank my advisor and the first auditor of this thesis, Silvia Miksch, for her great professional support but also for her great leadership skills as she always finds the right balance between guidance and giving me plenty of rope. Likewise, I want to thank the second auditor of this thesis, Bernhard Preim, for his time and expert advice. I also want to thank Katharina Kaiser for her most helpful support in the evaluation process of CareCruiser. Moreover, I am very grateful to Wolfgang Aigner for always providing professional feedback and guidance and also to the rest of the Information Engineering Group (Bilal Alsallakh, Paolo Federico, Tim Lammarsch, Alexander Rind, and Andreas Seyfang) for the many fruitful discussions of my work. Moreover, I also want to thank the whole group, but Silvia Miksch in particular, for the friendly and constructive atmosphere and the comfortable conditions for carrying out my research.

In addition, I wish to thank Christian Popow, Department of Child and Adolescent Psychiatry, Medical University of Vienna, and Harald Leitich, Department of Obstetrics and Fetal Medicine, Medical University of Vienna, for their cooperation, helpful suggestions, and carefully testing the CareCruiser prototype. For their time and help in evaluating the CareCruiser prototype, I also want to thank Bilal Alsallakh, Paolo Federico, Simone Kriglstein, and Sarah Wiener.

Most importantly, I want to thank my father, Helfried Gschwandtner. He never got tired of discussing my research with me, of encouraging me, pushing me, and extensively proofreading this thesis. I am most thankful for his great help, time, and motivation work, which made this work possible.

The research leading to this thesis has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 216134 and from the Centre for Visual Analytics Science and Technology CVASt (funded by the Austrian Federal Ministry of Economy, Family and Youth in the exceptional Laura Bassi Centres of Excellence initiative).

Abstract

In recent years, sophisticated visualization methods have been developed to support both the logical structure and the time-oriented aspects of computer-executable clinical treatment plans. However, visualizing the effects of applying treatment plans as well as supporting the exploration of effects on the patient's condition are still largely unresolved tasks.

To fill this gap, we have developed an interactive Information Visualization prototype (called CareCruiser) that enhances known visualization methods to communicate the processes of treatment plan applications together with their effects on a patient's condition in an intuitive way.

Our prototype combines the advantages of enhanced visual pattern recognition with the traditional information concerning the development of parameters. Thus, it provides means (1) to assess success or failure of previously applied treatment plans, (2) to explore the effects of each applied clinical action on the patient's condition, and (3) to identify sub-optimal treatment choices. These means help physicians to optimize their treatment choices and enable developers of clinical practice guidelines (CPGs) to investigate and readjust these treatment plans.

To ensure the quality of our Information Visualization prototype (1) we collaborated with a medical expert in an early design phase, (2) we conducted a heuristic usability evaluation, and (3) we tested its usefulness with domain experts. In the course of the usability evaluation the evaluators tested the prototype in conformance with generally accepted usability principles; in the user testing, physicians were asked to function as users and were interviewed afterwards. Both the usability evaluation and the user testing led to useful results and presented a very positive picture of the visualization and exploration features of CareCruiser.

Kurzfassung

Betrachtet man visuelle Darstellungen von (computerausführbaren) klinischen Therapieplänen, so wurden in den letzten Jahren einige bemerkenswerte Visualisierungen entwickelt um, einerseits, die interne Logik dieser Therapiepläne, und andererseits, deren zeitbezogenen Aspekte darzustellen. Jedoch gibt es bisher weder visuelle Darstellungen der Auswirkungen der angewandten Therapien auf den Patientenzustand noch gibt es Möglichkeiten diese Auswirkungen interaktiv zu untersuchen.

Ausgehend von bekannten Visualisierungsmethoden haben wir einen interaktiven Prototyp (namens CareCruiser) entwickelt, der die Anwendung von Therapieplänen und deren Auswirkungen auf den Patientenzustand auf leicht verständliche Weise kommuniziert und untersuchbar macht. Unser Prototyp kombiniert die Vorteile von visueller Mustererkennung mit gewohnten Liniendiagrammen, welche die Entwicklung von bestimmten Patientenparametern intuitiv darstellen. Dadurch haben wir ein Werkzeug entwickelt, um (1) den Erfolg von angewandten Therapieplänen zu beurteilen, (2) die Effekte der einzelnen klinischen Aktionen auf den Patientenzustand zu untersuchen, und (3) suboptimale Therapieentscheidungen zu identifizieren. Zum einen hilft das Ärzten ihre Therapieentscheidungen zu optimieren und zum anderen ermöglicht das den Entwicklern von klinischen Leitlinien diese in Bezug auf ihre genauen Auswirkungen zu untersuchen und gegebenenfalls zu korrigieren.

Desweiteren haben wir verschiedene Schritte unternommen um die Qualität unserer interaktiven Visualisierung zu bemessen: (1) in einer frühen Entwurfs-Phase haben wir mit einem Medizin-Experten zusammengearbeitet, (2) wir haben eine heuristische Usability-Evaluierung durchgeführt und (3) wir haben die Brauchbarkeit unserer Visualisierung mit den tatsächlichen Endbenutzern (Ärzten) getestet. Zur Usability-Evaluierung haben wir allgemein anerkannte Usability-Grundsätze angewandt; beim Test mit den Endbenutzern haben wir zwei Ärzte gebeten die Funktionen unseres Prototypen zu testen und sie daraufhin interviewet. Sowohl die Usability-Evaluierung als auch die Benutzer-Tests führten zu brauchbaren Resultaten und lassen darauf schließen, dass die entwickelten Visualisierungsmethoden und Werkzeuge zum interaktiven Umgang mit den Daten vielversprechend sind.

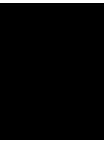
Contents

I	Analysis	1
1	Introduction	3
1.1	Motivation	3
1.2	Clinical Practice Guidelines and Protocols	5
1.3	Information Visualization	6
2	Problem Analysis	9
2.1	Medical Treatment Plans	10
2.2	Patient Data	12
2.3	Scientific Method	13
3	Related Work	19
3.1	Visualizing Hierarchical Data, Logical Sequences, and Time-Oriented Data . . .	19
3.2	Visualizing Patient Data	34
3.3	Visualizing Medical Treatment Plans	41
3.4	Discussion and Shortcomings of Existing Approaches	57
II	CareCruiser	61
4	Requirements	63
5	Design and Implementation	67
5.1	User Interface Composition	67
5.2	Visual Encoding	69
6	Comparing Effects and Discovering Temporal Patterns	83
6.1	Comparing the Conditions of Different Patients	84
6.2	The Success of the Applied Treatment Plan	84
6.3	Detecting Temporal Patterns in Reactions to Clinical Actions	87

III Evaluation, Results, and Transferability	93
7 Multi-Step Evaluation Process	95
7.1 Collaboration with a Medical Expert	95
7.2 Usability Evaluation	96
7.3 User Testing	97
8 Outcome	101
8.1 Usability Evaluation	101
8.2 User Testing	102
9 Transferring to Other Domains	105
IV Conclusion & Future Directions	109
9.1 Summary	111
9.2 Answers to Research Questions	112
9.3 Publications	114
9.4 Main Contributions	115
9.5 Further Work	116
V Appendix	119
A User Study Protocol	121
B Transcribed Interviews	125
B.1 Interview I	125
B.2 Interview II	129
Bibliography	133

Part I

Analysis



Introduction

*‘The eye...
the window of the soul,
is the principal means
by which the central sense
can most completely and
abundantly appreciate
the infinite works of nature.’*

Leonardo da Vinci (1452 - 1519)

1.1 Motivation

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 has become imperative. Card et al. defined InfoVis as ‘The use of computer-supported, interactive, visual representations of [abstract] data to amplify cognition’ [13, p.6]. By providing interactive visual representations, InfoVis aims at supporting intelligibility and improved exploitation 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’ [4, p.1]. Furthermore, it allows the user to explore data at different degrees of detail to avoid overload, e.g., by 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, namely visualizing objects of the physical world, are associated with the discipline of Scientific Visualization [76]. In contrast to this, InfoVis 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 (such as patients' data or treatment processes), the visualization of computerized clinical treatment plans is a challenge for InfoVis researchers which has not been targeted until quite recently [4].

From the plan management point of view, clinical treatment plans can be seen as (time-dependent) plans or activities [57]. In contrast to this, patients' data is composed of heterogeneous information spaces, including quantitative and qualitative data, high-frequency and low-frequency data, etc. Conceptually, patients' data can be seen as multidimensional information spaces. For both concepts a variety of visualization approaches have been developed in order to meet specific requirements [4]. For the visualization of a patient's condition in combination with the applied clinical treatment, though, only very few contributions [1,5,86] exist. 'Usually, the visual representations of clinical treatment plans are simply combined with text-based explanations of the course of the patient' [4, p.8]. However, visualizing this connection would be of great benefit, since changes of the patient's condition in accordance with the execution of the treatment plan are of major interest to physicians. There is no appropriate visualization method yet to illustrate these dependencies [4]. Thus, '...visualizing patients' data in connection with clinical treatment plans is still an open, but challenging issue' [4].

This work is aimed at filling this gap. It is an instrument to explore visually and interactively (1) in which way the application of specific therapy plans as well as the application of selected clinical actions cause changes in the health condition of a patient and (2) how well these treatment plans succeed in reaching the intended effects.

Setting

Before taking a closer look at the subject fields of clinical practice guidelines and Information Visualization, some words should be said about the bigger context in which this work was carried out.

The ReMINE project [74] aimed at facilitating an early identification and an effective prevention of Risks Against Patient Safety (RAPS). In the context of this project, my work was to design and develop InfoVis methods in order to display the patient data in combination with applied clinical treatments and to make explorable both the applied clinical actions and how they influence the patient's condition.

Moreover, my work was also carried out on behalf of the Laura Bassi Centre for Visual Analytics Science and Technology (CVASt) [21]. In the course of the CVASt project, we aim to develop novel, user-oriented, and task-specific Visual Analytics [91] methods. In particular, we focus on the development of tools that take advantage of the strength of both, the enormous processing power of computers and the humans' visual capabilities of information processing. In doing so, we aim to provide novel tools to foster new insights and enable knowledge discovery.

In carrying out the design, implementation, and evaluation of this work, the close collaboration and many discussions with my working team and my other colleagues were of invaluable help for me. This is why I will refer to this work as 'our work' in the remaining thesis.

Overview of this Thesis

Part I deals with the introduction to and analysis of the research problem. After giving a short introduction to the context of this thesis in Chapter 1, we analyze the problem at hand, our research questions, and discuss the methods we chose to tackle it in Chapter 2. In Chapter 3 we give a comprehensive description of the state-of-the-art of methods suited to visualize medical and time-oriented data. Subsequently, we discuss how well these methods are suited to solve our problem in Section 3.4.

Part II deals with the design and implementation of an original, interactive InfoVis prototype to meet our requirements (listed in Section 4). The following chapter (Chapter 5) describes the conceptual design and the implementation of our visualization prototype – called CareCruiser – in detail. In Chapter 6 we give three examples to illustrate how CareCruiser can be used in different clinical situations.

Part III is dedicated to the evaluation of our prototype and the transferability of its visualization methods: In Chapter 7 we describe the guided design process (gathering feedback from a medical expert), the usability evaluation, and the user study we have conducted. Subsequently, in Chapter 8, we discuss the results of these evaluations. In Chapter 9 we discuss how the visualization methods developed for CareCruiser could be applicable with advantage to other domains, and we give two examples with simulated data.

Finally, in Part IV, we summarize the steps that have been taken and the answers we can provide to our research questions (Section 9.2). We give an overview over publications of our findings at international conferences and in journals (Section 9.3) and we state our main contributions in Section 9.4. Finally, in Section 9.5, we discuss our plans of further work in this field.

The appendices include the user study protocol and the transcribed interviews of this study, followed by the bibliography.

1.2 Clinical Practice Guidelines and Protocols

This subsection deals with the definitions of the two terms ‘clinical practice guideline’ and ‘protocol’, which are essential terms in the context of this thesis.

1.2.1 Clinical Practice Guidelines

Clinical practice guidelines (CPGs) are aimed to assist practitioner decisions:

‘Clinical practice guidelines are systematically developed statements to assist practitioner and patient decisions about appropriate health care for specific clinical circumstances.’ [28, p.38]

CPGs provide recommendations on the appropriate treatment of specific conditions according to best available scientific evidence. Thus, they are important means for quality control, they facilitate clinical decision making, they establish standards to reduce inappropriate variation in health care, and they can be used for the instruction and training of physicians [25, 28, 94].

There are different ways to develop CPGs. Some CPGs are based on the opinion of experts in specific fields, others rely on a formal consensus process, and evidence-based guidelines rely on a systematic review of clinical evidence [94].

1.2.2 Protocols

Protocols are more specific than CPGs:

‘Protocols are local tools that set out specifically what should happen, when and by whom in the care process. They can be seen as the local definition of a particular care process derived from a more discretionary guideline. They are in essence tools that assist in quality improvement and reducing inequalities. ... Protocols reflect local circumstances, and variation will due to the differing types of local provision.’ [89]

Protocols differ from CPGs in that they are defined in greater detail and give precise instructions on what to do at specific institutions.

1.2.3 Computerized Guidelines

Being intended to improve the quality of patient care and to reduce costs, CPGs are increasingly being used. However, CPGs provide population-based recommendations while clinician decision-making is most effectively supported by providing patient-specific advice delivered during patient encounters [64]. The availability of patient-specific advice at the point-of-care turned out to be crucial for effectively using guidelines in clinical care [25]. To this end, making CPGs computer-executable paves the way to guideline-based point-of-care decision support systems, which entail numerous benefits [25, 94], for instance:

- They provide recommendations tailored to an individual patient;
- They don’t interrupt clinical workflow; instead, they provide integrated computer support;
- They reveal inconsistencies and errors in the text of the guideline;
- They provide reminders and alerts.

Therefore, several groups and organizations work at the development of computerized guidelines and decision support systems that incorporate these guidelines.

1.3 Information Visualization

Presumably the most widely accepted definition of Information Visualization is this:

‘The use of computer-supported, interactive, visual representations of [abstract] data to amplify cognition’ [13, p.6].

Within this definition, the term cognition refers to ‘the intellectual process in which information is obtained, transformed, stored, retrieved, and used’ [12, p.515]. This intellectual process could be, for instance, writing a scientific paper or shopping in Internet shops.

Card describes the clear benefit of visualization as using the world outside the mind as a resource for thought [12]. Interacting with the world outside is a great advantage for the working mind. Examples that strengthen this claim are grocery lists and pocket calculators. Both of them are very simple means of interacting with the external world which ease an otherwise more difficult task by taking away the need to remember every single detail. By using visualization as an external resource of thought they [13]

1. increase the memory and processing resources available to the users,
2. reduce search for information
3. use visual representations to enhance the detection of patterns,
4. enable perceptual inference operations,
5. use perceptual attention mechanisms for monitoring, and
6. encode information in a manipulable medium.

Moreover, the visual sense is the sense with the highest information capacity which makes visualization the means of choice to communicate huge amounts of complex and interrelated data in a feasible way.

However, Information Visualization needs to be distinguished from related areas:

- **Scientific Visualization** is focused on scientific data which is mostly physically based, i.e., the visualization of invisible information within an existing spatial framework (e.g., the human body or earth geography) [12].
- **Data Graphics** is ‘the use of abstract, nonrepresentational visual representations to amplify cognition’ [12, p.515]. Both Information Visualization and Scientific Visualization are sub-fields of Data Graphics.
- **Information Design** is the next-largest field which is defined as ‘external representations for amplifying cognition’ [12, p.515].
- Finally, **external cognition** is ‘the uses of the external world to accomplish some cognitive process’ [12, p.515].

In contrast to Scientific Visualization, Information Visualization is concerned with finding an effective mapping between abstract entities and visual entities to represent them as well as their spacial arrangement. An effective representation of these abstract entities aims at helping the user to gain insights visually.

Problem Analysis

In a first step, we conducted a comprehensive literature review to investigate the different user groups of a computerized CPG (computerized, medical treatment plans) as well as their specific information needs [39]. According to our findings physicians most frequently gathered information about the following categories (listed according to their priorities):

1. Information about treatment (e.g., recommended treatment, specific procedures, effects and side effects); of particular importance are information on medication (e.g., recommended medication, effects and side effects, efficiency).
2. Information relevant for diagnosis (e.g., possible causes of symptoms, relations between findings, indicated tests).
3. Information on patient data (e.g., symptoms, medical history, test results).
4. Information about other health care providers (e.g., contact data of colleagues and other specialties).
5. Information about a specific disease (e.g., prognosis, etiology).

When focusing on information categories that are relevant in the course of executing 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.

It is crucial to visualize the currently applied treatment plan, what therapy or medication is being applied at the moment, which treatments have been applied so far, a.s.o. Apart from 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, a.s.o. 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's condition has to be represented in connection with the applied treatment plans in a readily interpretable way.

2.1 Medical Treatment Plans

Asbru is a time-oriented, intention-based, skeletal plan-specification representation language [58, 78] (used in the Asgaard Project [81]) to express clinical protocols or CPGs. There are several other languages for executable or computer-interpretable guidelines (see [64]), but only Asbru and PROforma support expressing intended effects of actions [64]. Asbru defines intentions for guidelines which can be seen as goals that should be fulfilled. Thus, we decided for the Asbru language to develop an interactive InfoVis prototype to explore the effects of guidelines and clinical actions.

We had to analyze the type and structure of Asbru plans in order to find a suitable representation, as treatment plans formulated in Asbru show some complex characteristics:

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

2.1.1 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. 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 [78]. Thus, a plan and its sub-plans form a hierarchical structure, which is relevant for their execution.

2.1.2 Logical Sequences and Execution Order

In accordance with 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.

2.1.3 Complex Knowledge Roles

Preferences define the applicability of a plan, i.e., they serve as selection criteria as to how well the plan fits the patient’s condition (e.g., exact fit, approximate fit); they also 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 [58]:

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 last after finishing the plan (e.g., patient has less than one high glucose value per week);
4. Overall action pattern: the overall pattern of actions that should last after finishing the plan (e.g., patient has visited dietitian regularly for at least three months).

This information is important not only for selecting the right plan but also for scrutinizing treatment plans as part of the ever ongoing process of improving the treatment. This makes intentions one of the key parts of Asbru [78].

Conditions need to be complied 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 are complied), it has reached its goals, whereupon the next plan in the sequence is to be executed [78].

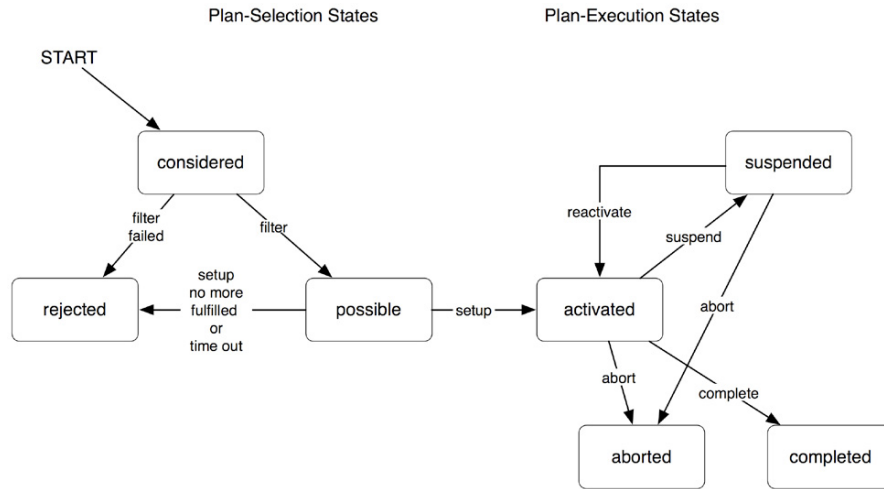


Figure 2.1: Different states of Asbru plans.

Effects describe the relationship between plan arguments and measurable parameters by means of mathematical functions. A probability of occurrence is also given [78].

2.2 Patient Data

Due to the use of modern information technologies in clinical care, enormous amounts of heterogeneous data and information are generated for a single patient. These kinds of information are often interrelated and complex. Patient parameters cannot be examined in isolation as there are always a couple of parameters that are closely correlated. Moreover, the parameter value at a single point in time is less meaningful than its evolution over time. Except for medical images such as x-ray images or magnetic resonance images, patient data are mainly made up of abstract data – often numerical values such as fever curves or blood parameters, but also qualitative types of data such as positive and negative findings. In summary, patient data [4]

- are enormous amounts of complex data,
- contain various parameters which are interrelated,
- change over time,
- are often abstract, and
- can be quantitative or qualitative.

Moreover, looking at a patient's condition in combination with the applied treatment is crucial to optimize treatment choices. In other words, the identification of changes in a patient's condition in reaction to applied treatments – i.e., the bettering or worsening of single patient parameters, the overall success or failure of applied treatments, individual patient reactions on specific clinical actions, etc. – deserves maximum attention.

2.3 Scientific Method

In a first step, we defined the problems we wanted to solve. Focusing on CPGs and how they can be improved, we formulated one main and four sub research questions.

2.3.1 Research Questions

Main question:

- Can we support a **detailed analysis of - and an adequate reaction to - the interdependencies** of a patient's health condition and the applied treatment (plan) with Information Visualization?

Sub questions:

- How can the **success or failure** of an applied treatment plan in terms of a **positive or negative change** in a patient's condition as a reaction to the applied treatment plan adequately be visualized?
- How can the visualization support the **comparison** of the effects of the **same treatment** on the conditions of **a number of patients**?
- Which means are suited to **detect temporal patterns** of the effects of clinical actions?

These questions were to serve both as an orientation guide through the project and as criteria to evaluate its outcome.

2.3.2 Literature Review and Analysis of Data Characteristics

We conducted a comprehensive literature review of information needs of medical personnel and decided to focus on information needs in context of applying treatment plans, i.e., information about treatments and patient data [39]. Subsequently, we conducted a detailed analysis and investigated the characteristics of treatment plans as well as the characteristics of patient data.

However, our main focus lay and still lies on the identification of effects of applied treatment plans (not only patient data nor treatment plans in isolation). These effects may range from well known effects over effects that were somehow expected to previously unknown or even surprising effects. Moreover, these effects may happen in general or as an individual patient reaction only. Thus, it would be hard to formulate a question before looking at the data. In addition, the number of events of interest (e.g., the effects of the application of a specific clinical action) in the given dataset may be too small to warrant the application of statistical methods. To this end, the visualization of the dataset serves to generate hypotheses which – of course – have still to be approved by a medical expert.

2.3.3 Decision for InfoVis Techniques

With respect to the findings of the analysis of the problem domain, we decided for InfoVis techniques.

Fekete et al. describe the characteristics of data and tasks to which InfoVis techniques are best applied [26]. According to them InfoVis techniques are of great use for exploratory tasks (e.g., browsing large amount of information) – sometimes without a specific goal or question in mind. These techniques enable the user to learn more about the data by examining them and thus, to [26]

- gain new insights,
- find unexpected behavior,
- formulate better, more meaningful questions and hypotheses, and
- gain a broad understanding of an information space.

Moreover, Lin describes a number of conditions in which browsing may be especially useful [55]. Among others Lin lists the condition ‘When information is easier to recognize than describe’ [55].

According to the benefits of InfoVis described above, sophisticated InfoVis techniques are the means of choice to deal with the available information about patient data and the applied treatment in a comprehensible way, browse the data, and potentially formulate hypotheses about the effects of specific clinical actions.

2.3.3.1 Information Visualization vs. Automatic Analysis

There are other well grounded scientific methods suited to understand complex data, for instance Statistics and Data Mining which is also concerned with automatically understanding data structures and patterns within the data. Fekete et al. outline what InfoVis can do that statistics and data mining cannot [26].

Statistics is composed of several subfields such as classical statistics, descriptive statistics, Bayesian statistics, and Exploratory Data Analysis. InfoVis can be seen as an extension of Exploratory Data Analysis. To solve a given problem by making sense of the gathered data, classical analysis starts by designing a model of the data and subsequently applies mathematical analysis to test the model. Usually Exploratory Data Analysis also tries to find a model of the data but it uses visual exploration methods to acquire insights as to what the data looks like. Anscombe [8] provides four datasets with the same statistical profile, namely

- same mean of the x values,
- same mean of the y values,
- same equation of the least-squared regression line,

- same sums of squared errors,
- same regression sums of squared errors,
- same residual sums of squared errors,
- same correlation coefficient, and
- same coefficient of determination.

However, the actual shapes of these four datasets are quite different from one another (see Figure 2.2).

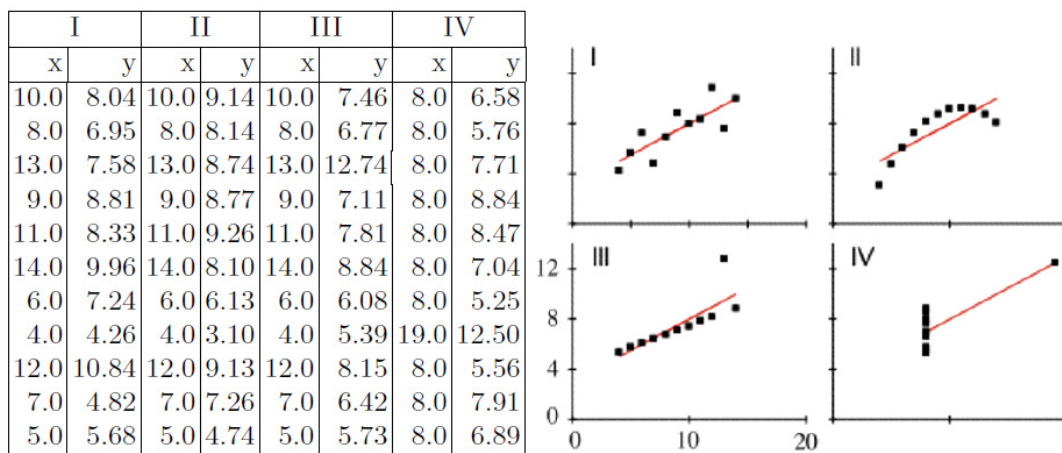


Figure 2.2: Four datasets by Anscombe [8] with same statistical profile but quite different shapes [26].

Consequently, in case we don't know much about the nature of the data it is useful to visualize the data before trying to find a model [26].

Data Mining is even more concerned with automatically finding interesting facts in large datasets. Spence and Garrison [85] address the question whether Data Mining can automatically solve these kinds of problems instead of humans having to visually explore the dataset. As a refuting example they name the Hertzsprung Russell Diagram (see Figure 2.3 on the left side). This diagram shows the relation between the temperature of stars and their magnitude (= brightness) in a simple plot. Humans have interpreted this diagram as shown in Figure 2.3 (on the right side). However, no automatic method could compute this interpretation of the plot, due to the noise and artifacts on the data. The human vision system is still outstanding when it comes to filtering-out noise from data and spotting interesting patterns, which, by the way, is an important ability for hunters in deep forests [26].

To put it in a nutshell, the different methods (i.e., automated and exploratory methods) answer different questions.

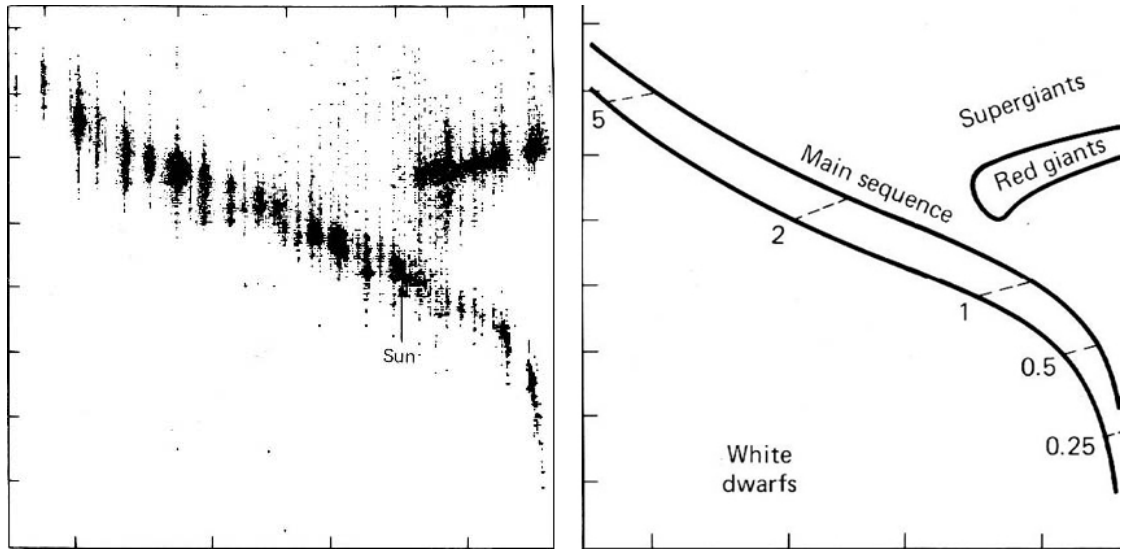


Figure 2.3: The left side shows the Hertzsprung Russell Diagram and the right side shows its standard interpretation by humans [85]. The X-axis represents the temperature of stars and the Y-axis represents their magnitude.

‘When a model is known in advance or expected, using statistics is the right method. When a dataset becomes too large to be visualized directly, automating some analysis is required. When exploring a dataset in search of insights, Information Visualization should be used, possibly in conjunction with Data Mining techniques if the dataset is too large.’ [26, p.14].

2.3.4 Literature Review of Related Work

After deciding for InfoVis techniques, we investigated to which degree existing InfoVis approaches meet the given requirements. Thus, we conducted a comprehensive literature research regarding *design*, *implementation*, and *evaluation* of visualizations of *patient data* and *clinical practice guidelines*. We searched the archives of topic related Journals (Artificial Intelligence in Medicine, International Journal of Medical Informatics, Journal of the American Medical Informatics Association, Information Visualization, etc.) as well as scientific, electronic databases (PubMed, ACM Digital Library, IEEE Xplore, and also Google Scholar) and proceedings of relevant conferences (International Conference on Information Visualization, IEEE Conference on Visual Analytics, ACM SIGCHI Conference on Human Factors in Computing Systems, etc.). Our search terms included ‘visualization of effects’, ‘visualization of patient data’, ‘visualization of guidelines’, and variations of these phrases. Moreover, we scanned the reference lists of relevant papers.

We looked both into earlier work, such as Petri Nets [66] (see Section 3.1.2.3) which are designed to visualize logical sequences (relevant in the context of treatment plans), and into more recent work, such as LiveRAC [56] by McLachlan et al. (see Section 3.1.3.7), an interactive visualization of system management time-series data (both treatment plans and patient data are

highly time-oriented) and PatternFinder [68] by Plaisant et al., which allows for the analysis of temporal patterns in patient histories (see Section 3.1.3.3); all of which are state-of-the-art research papers.

2.3.5 Prototypical Implementation

Based on these findings, we have designed and prototypically implemented an interactive visualization. As selected visualization approaches already provide sophisticated solutions for certain aspects – just to name one prominent example, CareVis [5] handles the visual representation of complex information of treatment plans (i.e., temporal constraints, nesting of plans, and plan characteristics) in an intelligible and comprehensive way –, we decided to take these approaches as a basis to develop a more consistent visualization. In addition to an appropriate visualization of patient data and treatment plans, easily usable interactive means had to be provided to allow the user to explore the data, thus enabling him/her to derive decisive insights into the effects of different therapies.

We went for a user-centered design process by creating mockups of the InfoVis prototype and presenting them to a physician who's task it was to detect problems of understanding, utility, and usability at an early stage of development. We then modified the visualization prototype with respect to the outcome of this session.

2.3.6 Evaluation of the Prototype

There are two fundamentally different approaches in evaluating software prototypes: quantitative evaluation and qualitative evaluation.

Quantitative evaluation methods are methodologies in which precision is high and in which some declaration can be made about the possible generalization to a larger population [14].

'The experiments or studies involve a rigorous process of hypothesis development, identification and control of the independent variables, observation and measurement of the dependent variables, and application of statistics which enable the declaration of the confidence with which the results can be taken. [...] Once the changes to the dependent variables have been measured, statistical methods can be applied to understand the relative importance of the results.' [14, p.26] (see Figure 2.4 for a brief overview).

Usually these studies focus on a very specific quantifiable aspect of the software - such as the time needed to solve a task or how many errors the user did make. However, there are some specialties to InfoVis prototypes that require working towards a larger understanding of the benefits and shortcomings of a system. According to the definition of InfoVis, 'to amplify cognition' is the ultimate purpose of InfoVis [13]. This task lacks a ground truth in terms of objective and quantifiable metrics which makes it very hard to precisely assess the value of developed InfoVis techniques with clarity and accuracy [26]. Moreover, InfoVis techniques are frequently used to explore datasets. This exploratory process itself is likely to influence the way of using the tool (i.e., new questions and tasks may arise), which makes it more difficult to evaluate the system in terms of solving specific tasks [26]. This is why quantitative evaluation methods do not succeed in appraising the actual benefit of these interactive visualizations.

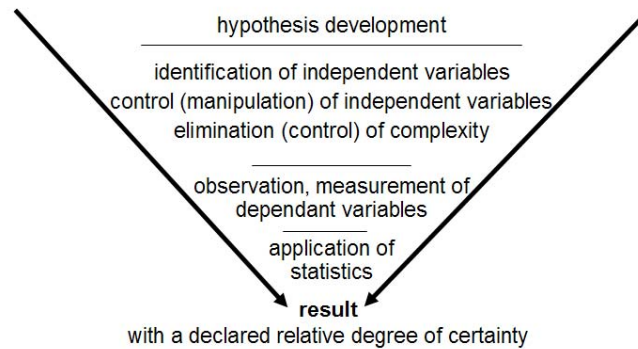


Figure 2.4: Schematic of the traditional (quantitative) experimental process [14].

Qualitative evaluation methods don't measure precise quantitative results but are aimed at achieving a richer understanding of using and interacting with the system and interacting with it in more realistic settings. To this end, the two primary methods for gathering data are observations and interviews [14]. Examples of qualitative evaluation methods are experimenter observations, think-aloud protocol, collecting participant opinions (i.e., questionnaire or interviews), and inspection evaluation methods such as usability heuristics [14].

Inspection evaluation methods focus on the usability of the interface. Since many of the tasks that are addressed by InfoVis techniques are interface interaction tasks, such as zooming, filtering, and accessing data details, usability is a key factor for InfoVis prototypes [84]. That is why we conducted a heuristic usability evaluation for which the participation of 3-5 usability experts is recommended [45].)

Moreover, in a late stage of the project, we conducted a qualitative assessment, i.e., a user testing involving two physicians, in the course of which we applied the think-aloud protocol [54] in combination with subsequent semi-structured interviews. The think-aloud technique involves encouraging the participants to phrase their thoughts while interacting with the prototype. It is a commonly used technique which enables the observer to gain some insight not only about how the domain expert works with the system but also about his/her thoughts, plans, and frustrations, which often out-weights the disadvantage that the participant may not feel familiar with this technique [14].

Related Work

3.1 Visualizing Hierarchical Data, Logical Sequences, and Time-Oriented Data

With respect to the findings of the analysis of the visualization-relevant characteristics of the data at hand, let us now take a closer look at existing approaches that deal with some of these special data characteristics. After looking at the different approaches in detail, we discuss their benefits and shortcomings in Section 3.4.

3.1.1 Visualizing Hierarchical Data

The nesting of treatment plans and sub-plans results in a hierarchical structure which has to be represented. The most common approaches of representing hierarchical relationships are trees, but there are also some more sophisticated approaches.

3.1.1.1 Trees

A tree is a visualization technique which represents the hierarchy by a directed and acyclic graph with one root node (see Figure 3.1(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; those without child nodes are called leaves. Usually these graphs are aligned from top to bottom (starting with the root node) representing each level of the hierarchy by means of a separate row.

Other approaches such as hyperbolic trees [53] place the root node in the center and align child nodes on circular layers around the root node (see Figure 3.2). 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 & 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. Other examples of representing trees are icicle trees (see Figure 3.3(a)) and sunburst trees (see Figure 3.3(b)).

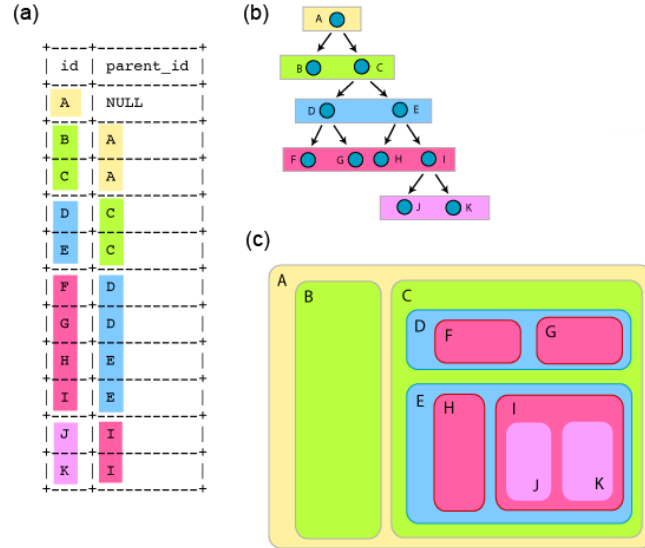


Figure 3.1: Most common ways to visualize hierarchical data: (a) listing all elements together with their parent elements, (b) a tree-layout starting with the root element A from which arrows lead to the child elements B and C, a.s.o., and (c) a treemap which uses 2D space to encapsulate child nodes into the area of their parent node [63].

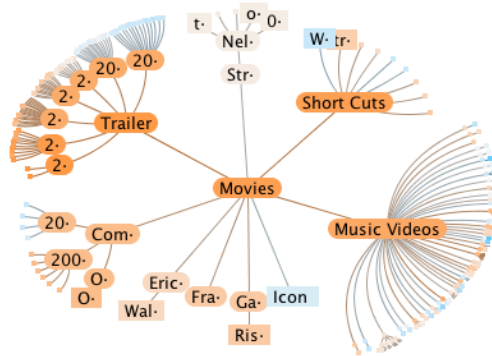


Figure 3.2: Hyperbolic Tree: the root node is located in the center and child nodes are aligned in circular layers around the root node. Interactive Focus & Context techniques allow for navigating the graph and investigating each node/area in detail while providing less detailed context information [73].

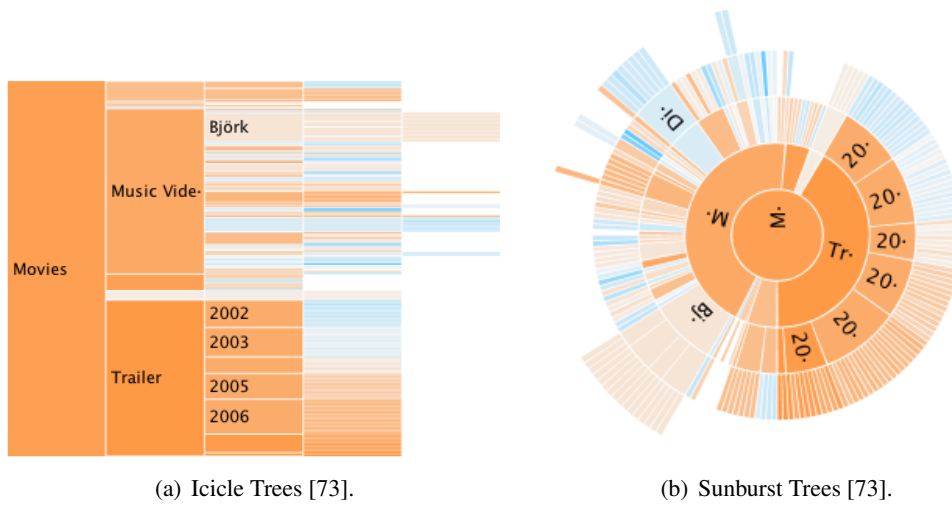


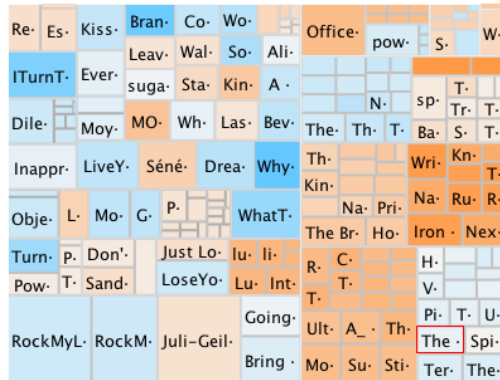
Figure 3.3: Additional space-filling visualizations of trees. Icicle Trees are space-filling representations which show the tree structure from left to right (sometimes also from top to bottom); rather than drawing links between nodes, the nodes are represented by areas and their placement indicates their position in the hierarchy. The size of the nodes is used to encode an additional quantitative dimension such as file size. Sunburst Trees are radial versions of Icicle Trees [73].

3.1.1.2 Treemaps

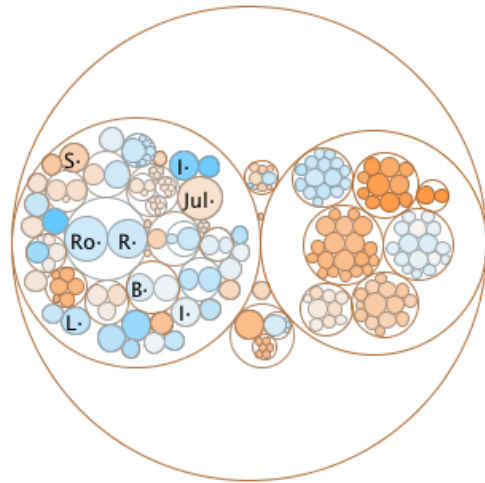
A treemap [83] is a rectangular, space-filling approach for visualizing hierarchical data (see Figure 3.1(c) and Figure 3.4(a)). It uses 2D space to lay out the area of the root element, which encapsulates the child elements. The size of each node is determined in relation to all other nodes; it represents a given attribute of the node (e.g., file size in the 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 a way to visualize different attributes of nodes and enable the user to compare nodes and sub-trees on different levels of the hierarchy. Typically, treemaps are rectangular, however, there are also circular treemaps (see Figure 3.4(b)). One disadvantage of circular treemaps is that they do not use the available display space as effective as rectangular treemaps do. On the other hand, the ‘wasted space’ of circular treemaps reveals the hierarchical structure better than rectangular treemaps [42].

3.1.1.3 Timeline Trees

Timeline Trees [11] address the challenge of visualizing hierarchical structures over time (see Figure 3.5). They are developed 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 an intermediate node the sub-tree starting at this node is collapsed; 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.



(a) Rectangular Treemap [73].



(b) Circular Treemap [73].

Figure 3.4: Treemaps: traditional treemaps are rectangular. Circular treemaps do not use the display space as effective as rectangular treemaps, but the hierarchical structure of the data is more visible in this way.

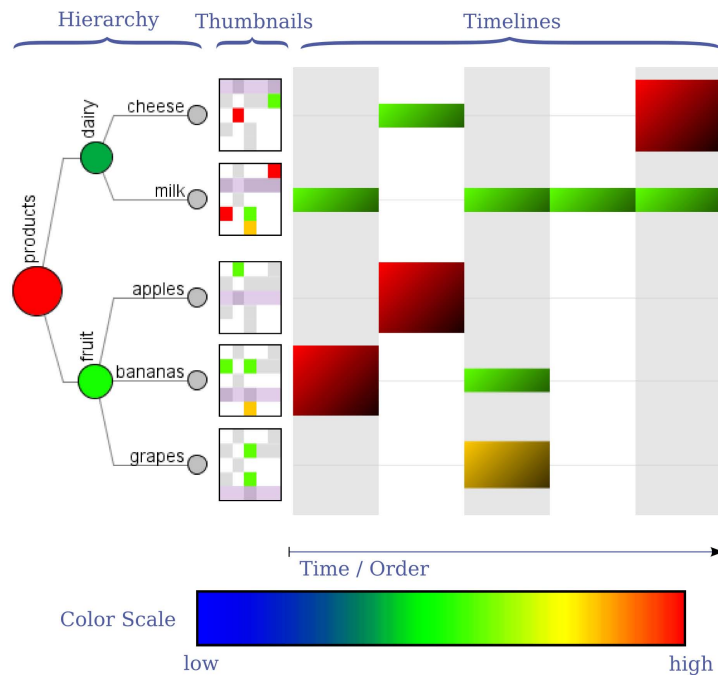


Figure 3.5: Timeline Tree: the hierarchy visualization on the left hand side is interactive and allows for expanding and collapsing nodes to select different layers of the hierarchy. The selection is propagated to the two other views, the thumbnails and the timelines [11].

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 of different size and color to encode 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.

3.1.2 Visualizing Logical Sequences

Above all, a visualization of clinical treatment plans has to communicate the logical sequence of actions. Existing approaches that deal with the visualization of logical sequences range from well-known rather basic diagrams to more sophisticated techniques, which are outlined in the following subsections.

3.1.2.1 Flow Charts

Flow Charts have been developed in 1947 [34] as a diagramming technique for the schematic representation of the flow of processes (see Figure 3.6). Initially Flow Charts were designed for software engineering, but they can be used to graphically represent any kinds of algorithms, for instance clinical algorithms.

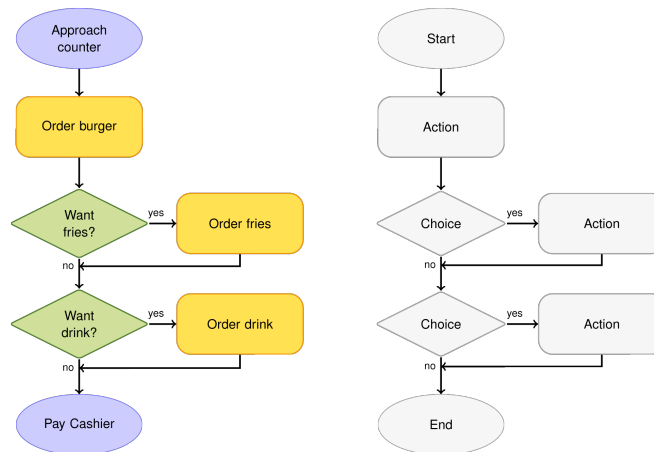


Figure 3.6: Flow Chart [65].

3.1.2.2 Nassi-Shneiderman Diagrams (Structograms)

In 1973 Nassi and Shneiderman developed a new method for illustrating algorithms and program functions [60]. With the help of this method the logical structure of a program can be laid down in detail. By enforcing a more structured program design the Structograms or Nassi-Shneiderman diagrams avoid some problems present in the very flexible Flow Chart diagrams. Moreover, they provide a very compact visualization of complex algorithms. Structural elements of algorithms are reduced to the basics: processes, sequences, branches, and loops, each

illustrated by specially shaped boxes (see Figure 3.7). Thus, the structural information of the algorithm can be learned from diagram layout (see Figure 3.8).

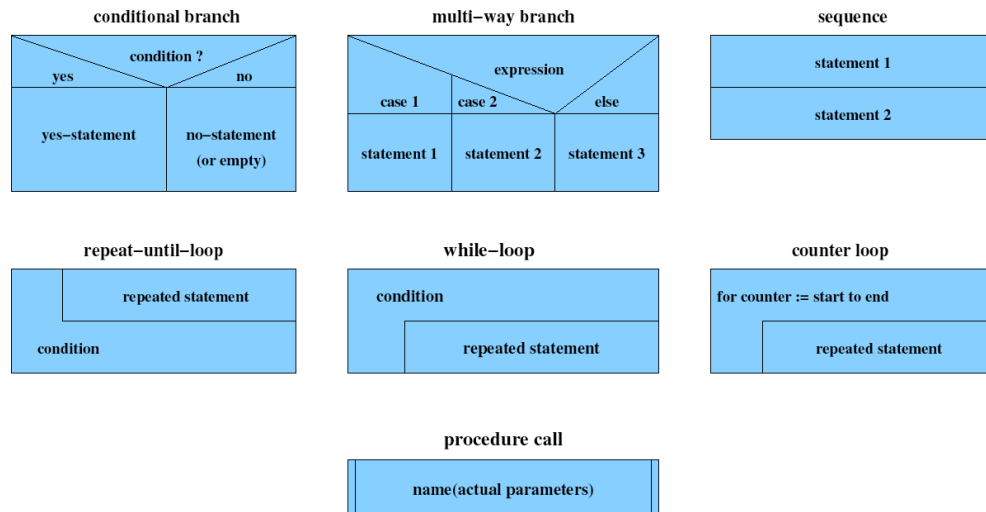


Figure 3.7: Structural elements of Nassi-Shneiderman diagrams [2].

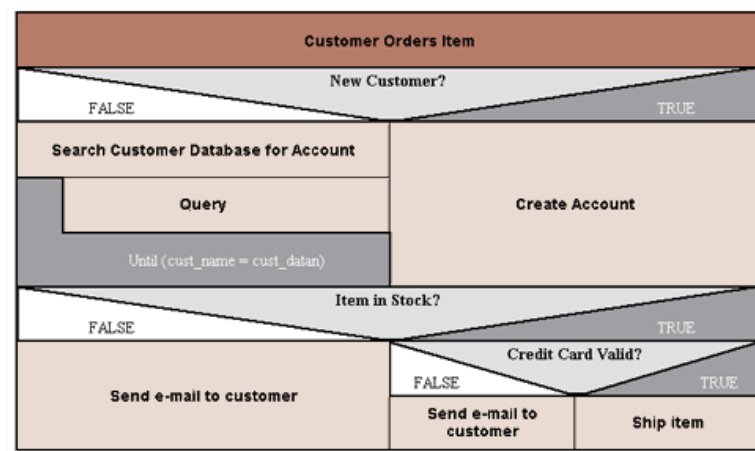


Figure 3.8: Example of a Nassi-Shneiderman diagram [49].

3.1.2.3 Petri Nets and State Transition Diagrams

Petri nets [66] and state transition diagrams [87] are directed graphs which describe the behavior of a system. A Petri net is a diagram that describes the concurrent processes which arise in a system with many components (distributed systems). It is composed of states (round symbols) and

transitions (rectangular symbols) (see Figure 3.9) [67]. In state transition diagrams, the nodes also represent states, while transitions are represented by labeled arrows (see Figure 3.10) [97].

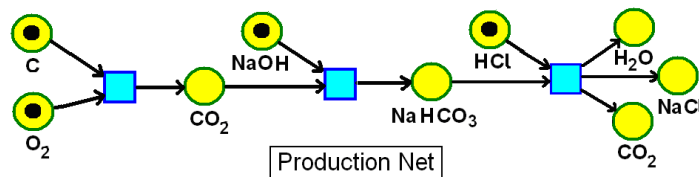


Figure 3.9: Example of a Petri net describing a chemical process [67].

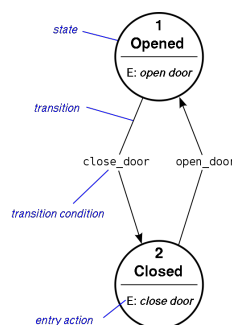


Figure 3.10: Example of a state transition diagram describing a door that can be opened and closed [97].

The main disadvantages of these visualizations (i.e., Nassi-Shneiderman diagrams, Petri nets, and state transition diagrams) are that they are not well known outside the software development community, and it takes some effort to learn the different diagram elements.

3.1.3 Visualizing Time-Oriented Data

For both treatment plans and patient data, the visualization of their temporal characteristics (e.g., start and end of the application of one treatment plan or the changes in the values of a patient parameter) are of major importance. There are several kinds of visualizations dealing with the representation of such temporal data. Due to space constraints, we outline only the most relevant temporal visualizations. For a comprehensive overview over visualizations of time-oriented data, however, we refer to Aigner et al. [6].

3.1.3.1 GANTT Charts

GANTT Charts were developed by Henry Laurence Gantt around the years 1910-1915 and are well known for project scheduling. They use bars on a timeline to illustrate phases and activities within their temporal context (i.e., the start, the end, and the duration of these activities; sequential, parallel, and overlapping activities). Sometimes they also indicate the dependency relationships between activities within the project (see Figure 3.11).

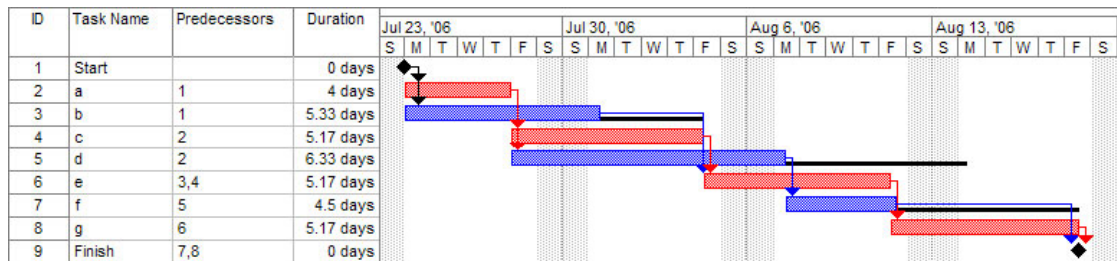


Figure 3.11: A GANTT Chart showing the time-oriented characteristics of phases and activities of a project as well as their dependencies [96].

3.1.3.2 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 (similar to GANTT Charts). This method is called Timeline; they are horizontal lines or bars depicting time spans and sequences and they are a powerful visualization method for time-oriented data (see Figure 3.12).

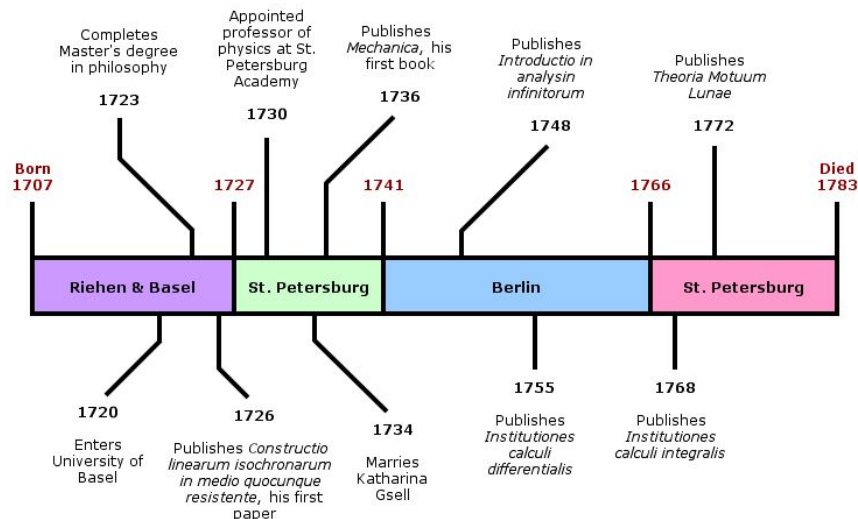


Figure 3.12: A Timeline of the life of Euler [90].

LifeLines [69, 70] is a research prototype for visualizing the time-oriented data of personal history records on a zoom-able timeline (see Figure 3.13). They use Time Lines with some extensions: among others, so-called ‘facets’ are introduced for grouping aspects of the record (in case of medicine, problems, allergies, diagnosis, labs, imaging, medications, immunizations, etc.) which can be expanded and collapsed. LifeLines was applied to summarize medical records as a set of lines and events. Moreover, LifeLines uses the color and height of individual bars to encode information: the fact that a bar is red may indicate a severe allergy, a status marked as

'alert' a.s.o.; the height of a bar may correlate to the significance of an action or event.

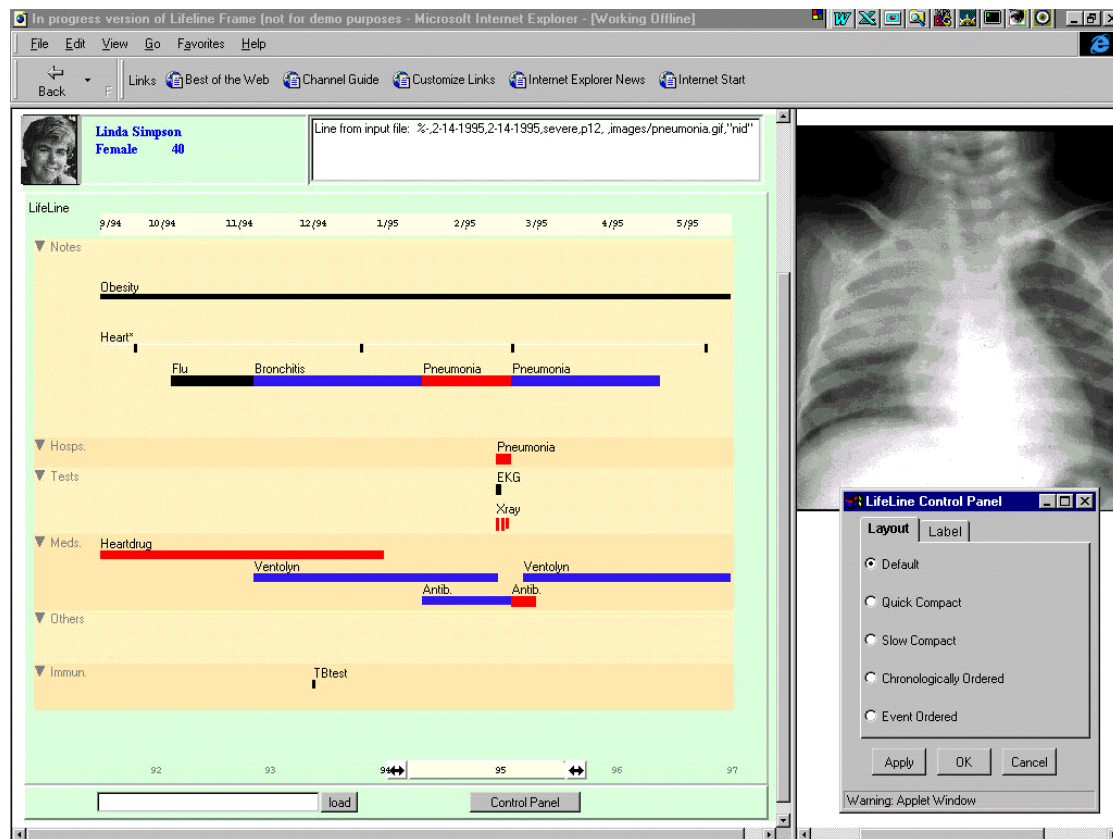


Figure 3.13: A patient record visualized by LifeLines; the chart is zoomed to reveal details of a long pneumonia episode [70].

LifeLines2 [95] is an extension of LifeLines and 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 previously unseen patterns (see Figure 3.14).

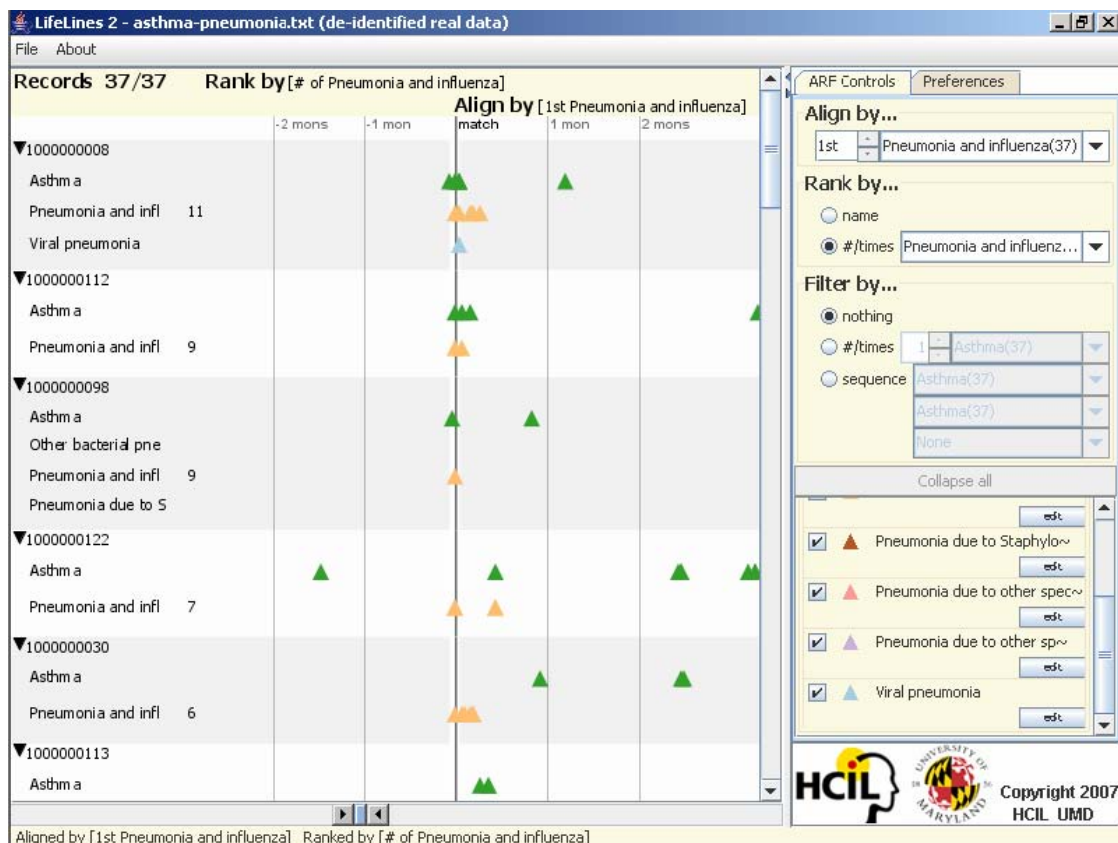


Figure 3.14: LifeLines2: All patient records are aligned by the 1st Pneumonia and influenza along a relative time scale. This alignment allows for an efficient identification of the diagnosis of asthma for the first 3 patients within a month prior or at the same time their pneumonia was diagnosed [95].

3.1.3.3 PatternFinder

The PatternFinder query tool [68] has been developed to facilitate the specification of queries in SQL and to subsequently visualize the results of the query. It has been implemented in a large operational system at the Washington Hospital Center. By providing means to easily specify SQL queries (e.g., '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.' [68]) and by providing a visual representation of the results, PatternFinder allows for the analysis of temporal patterns in patient histories. The user interface is divided into three panes (see Figure 3.15): the filter pane at the top left contains widgets to specify the query (i.e., the temporal scope, the sentinel event, the baseline event, the follow-on event, etc.). The results of the query are provided in a table in the bottom left pane; in addition, the pane at the right displays a visual representation of the results. A LifeLine [70] is drawn for each patient found; all events

that match the query are indicated. The user can expand the LifeLines chart for more detailed information about the events.

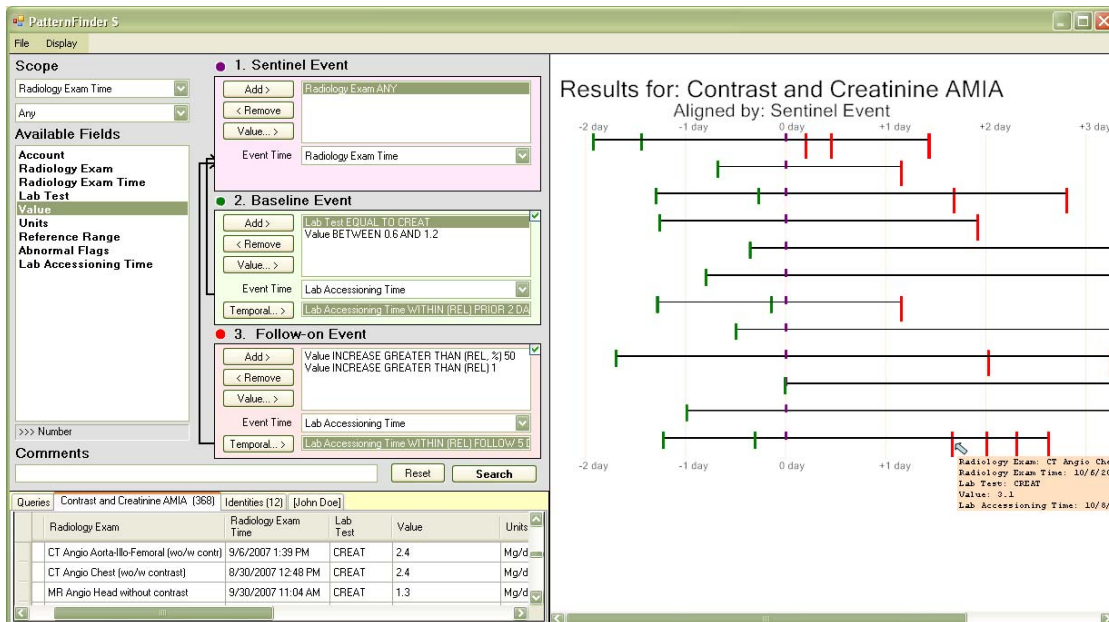


Figure 3.15: PatternFinder provides means to define a query in SQL, for instance: ‘Find patients who had a normal serum creatinine lab test less than 2 days before a radiology test with intravenous contrast, followed by an increase in serum creatinine by more than 50% and of more than 1.0 mg/dl within 5 days after the contrast administration’ [68]. For each patient found, a LifeLine with marks for matching events is drawn [68].

3.1.3.4 PERT Charts

Program Evolution and Review Technique (PERT) charts have been developed by the US Navy in 1950 to facilitate the planning and management of complex projects. PERT Charts graphically depict the single tasks of a project together with their temporal intertask dependencies by using boxes to represent tasks and arrows to depict predecessors and successors of each task. Scheduling information, such as

- the early start time (ES),
- the late start time (LS),
- the early finish time (EF),
- the late finish time (LF),
- the duration of the task,

- the buffer time (slack)

of each tasks are given as numerical values within the boxes (see Figure 3.16).

ES	Duration	EF
Task Name		
LS	Slack	LF

Figure 3.16: PERT Chart Element.

These charts are especially useful when the exact duration or start and end times of tasks are hard to define. Often PERT Charts are used to identify critical time constraints (the critical path) within the project, i.e., if there is no buffer time between the single tasks of the critical path; as a consequence, the delay of one of these tasks would always lead to a delay of the whole project.

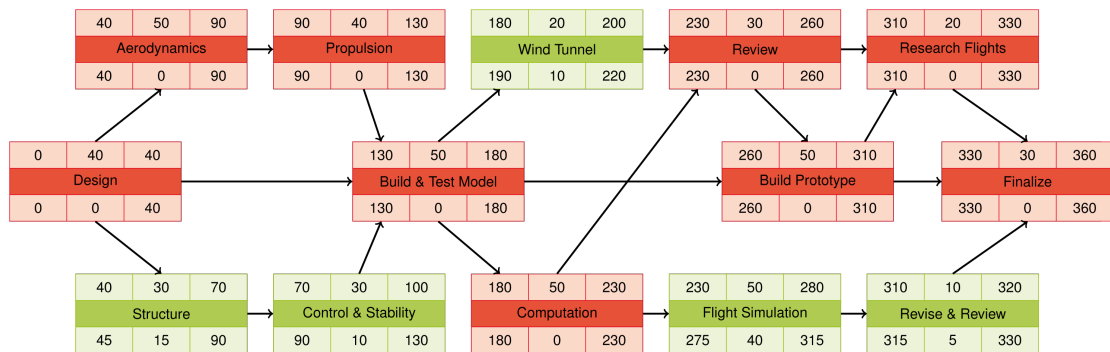
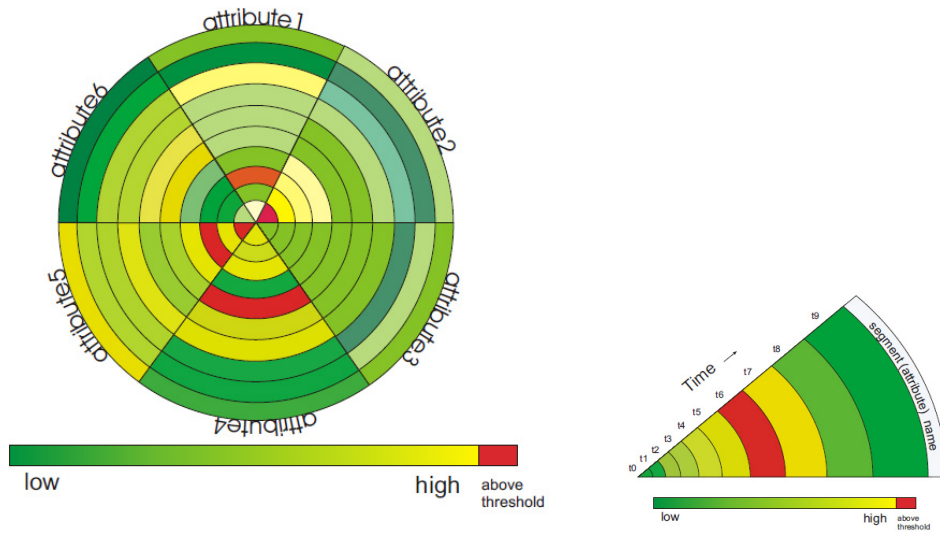


Figure 3.17: PERT Chart Example. The critical path is indicated by red boxes.

3.1.3.5 CircleView

Keim et al. [48] presented an approach for visualizing multidimensional time-oriented data sets. CircleView combines hierarchical visualization techniques and circular layout techniques with the intention to facilitate the comparison of continuous data changing their characteristics over time. It is aimed at detecting patterns, exceptions, and similarities in the data. Each segment of the CircleView represents one attribute (see Figure 3.18(a)). These segments are again subdivided into time-slot-segments giving a color-coded overview of changes of this attribute over time (see Figure 3.18(b)).



(a) Showing multiple attributes and their evolution over time [48]. (b) The evolution of time from the center to the edge [48].

Figure 3.18: CircleView.

3.1.3.6 TimeSearcher

Hochheiser and Shneiderman [44] developed a visual query tool for time series data, called Timeboxes, and integrated it into their software prototype TimeSearcher (see Figure 3.20). Timeboxes are visual, direct-manipulation queries for the interactive exploration of time series. Moreover, Hochheiser and Shneiderman extended their work to Variable Time Timeboxes (VTT), which allows for the specification of queries containing some uncertainty in the time axis [44].

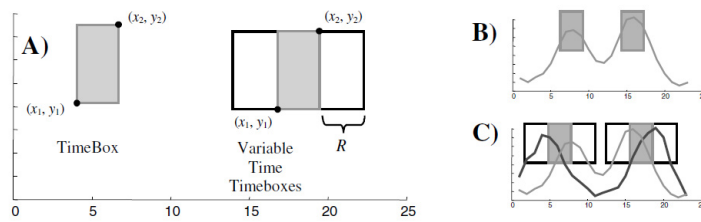


Figure 3.19: Variable Time Timeboxes (VTT). (A) Contrasting normal Timeboxes with Variable Time Timeboxes (VTT). (B) Normal Timeboxes qualify sequences that pass the shaded rectangle. (C) VTTS allow for the detection of patterns with a degree of uncertainty in the time axis [44].

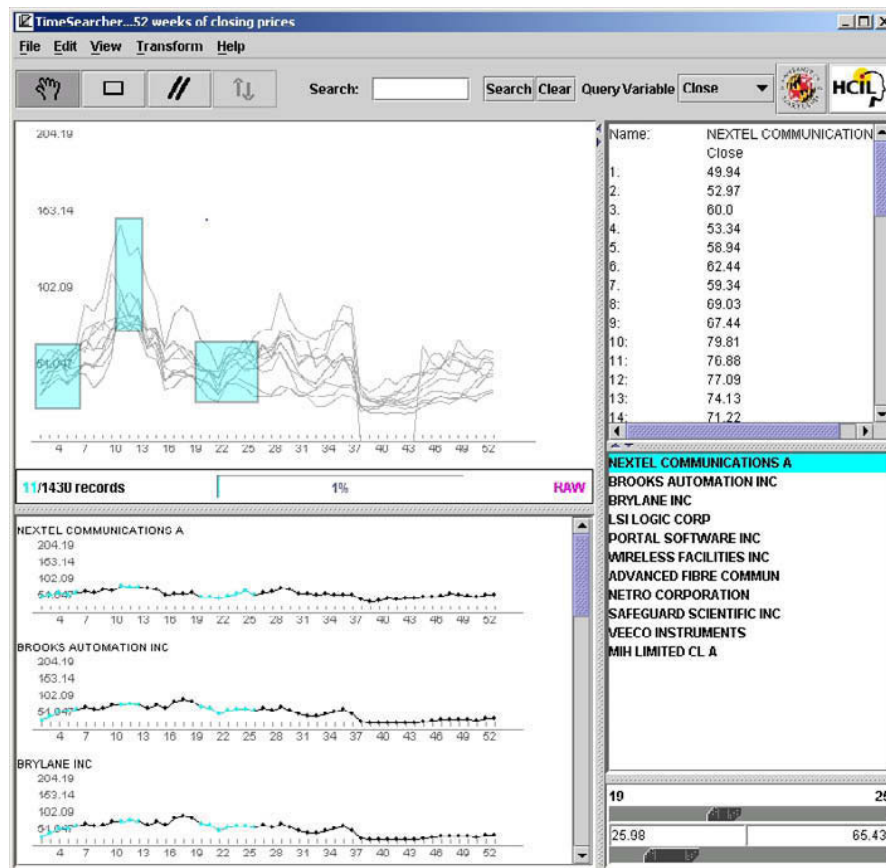


Figure 3.20: TimeSearcher. The pane in the upper left displays the visual query space; the pane in the lower left shows the queried data items [44].

3.1.3.7 LiveRac

One sophisticated approach of visualizing large collections of time series data is LiveRAC [56]. It arranges a huge number of charts in a matrix and provides features such as interactive re-ordering of rows and columns and Semantic Zooming which allows for exploring these charts at multiple levels of detail. A discrete color scale is used to color the cells' background, communicating specific, dynamically changeable values computed for each cell (i.e., thresholds of the minimum, maximum, or average values of the parameters)(see Figure 3.21). In addition, dots indicate the maximum value over the displayed time period (red dot), the minimum value (blue dot), and the current value (green dot).

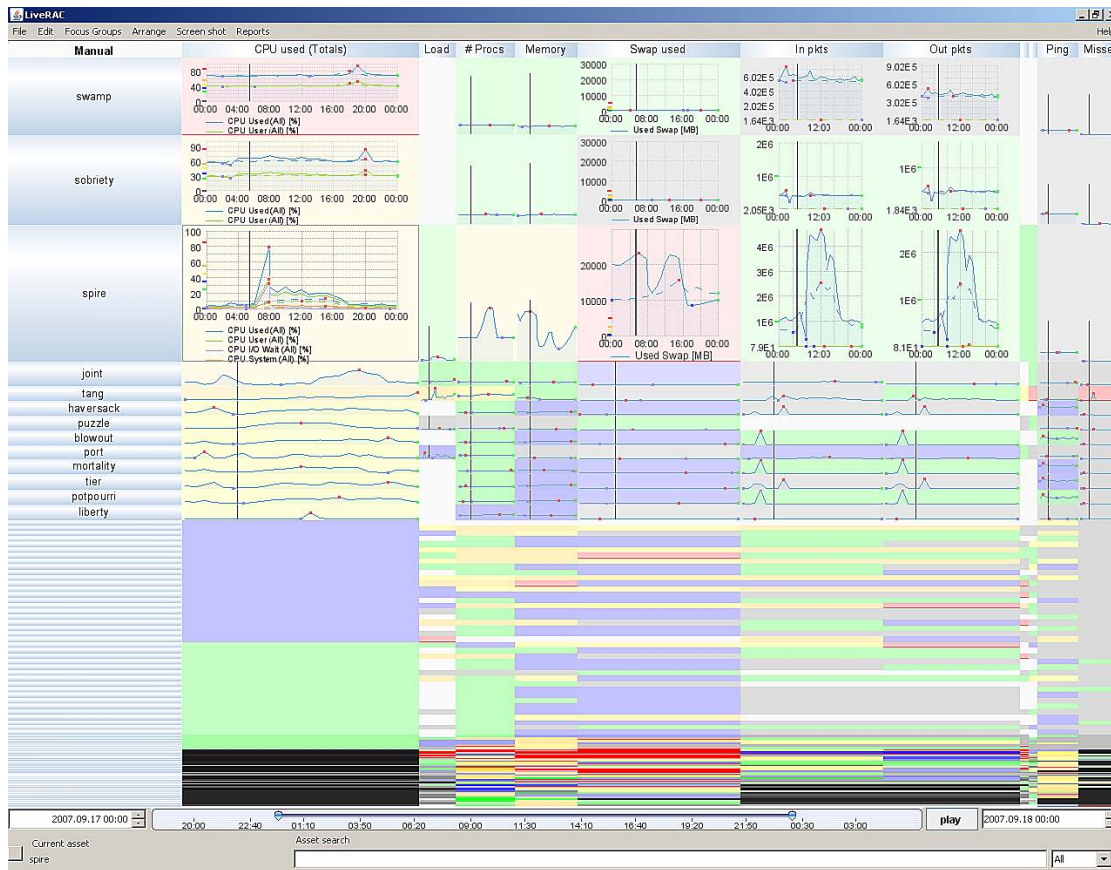


Figure 3.21: LiveRAC displays the data of a full day of system management time-series within a reorderable matrix. The rows can be individually enlarged to view the data at different Semantic Zoom levels [56].

3.1.3.8 Line Graph Explorer

The Line Graph Explorer [50] uses different colors, color saturation, and luminance to reduce the Y-dimension of line charts to 1D ribbons in order to visualize large collections of line graph data at limited display space.

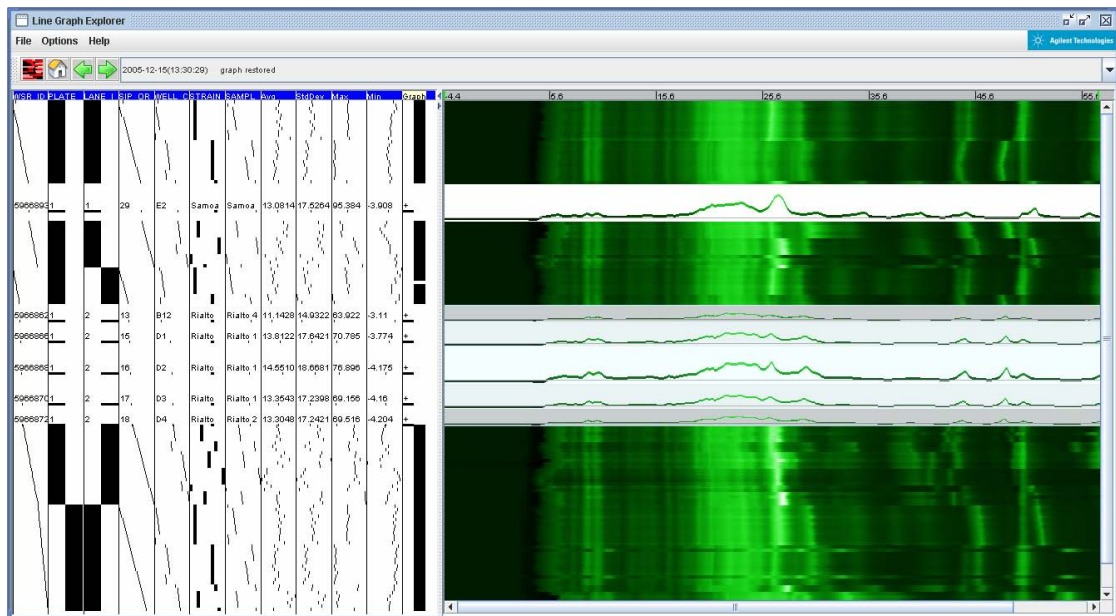


Figure 3.22: The Line Graph Explorer provides two main panels: the metadata panel on the left (which contains attributes of the line graph data) and the line graph panel on the right. Line-graphs can be viewed individually, in groups with a lens distortion, or an overview of multiple linegraphs given by color saturation and luminance [50].

3.2 Visualizing Patient Data

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

3.2.1 Graphical Summary of Patient Status

The Graphical Summary of Patient Status [71] uses small repeated graphs with identical formats to map different findings and treatments over time. The data of a 5cm thick record can be graphically summarized on a single page in this way. A non-linear time-scale compresses years of data into a context for assessing recent trends. Scaled values are plotted on a vertical axis where five same sized qualitative ranges (i.e., normal, elevated, reduced, critically elevated, and critically reduced values) are indicated (see Figure 3.23 and Figure 3.24). The Graphical Summary of Patient Status makes use of concepts like Small Multiples, Focus & Context, and the integration of textual and graphical information.

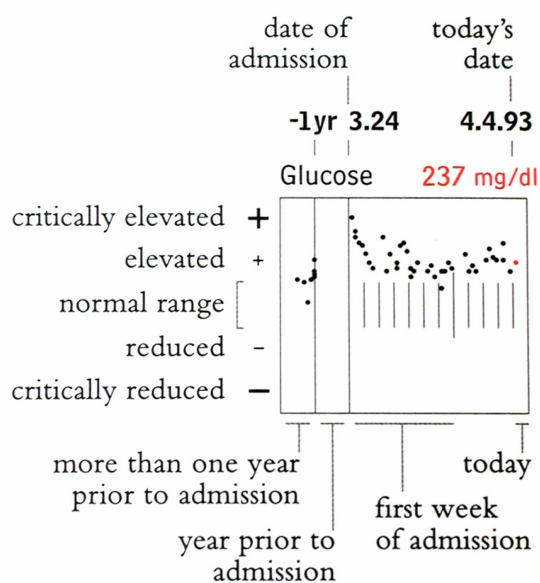


Figure 3.23: Graphical Summary of Patient Status – graph [71]

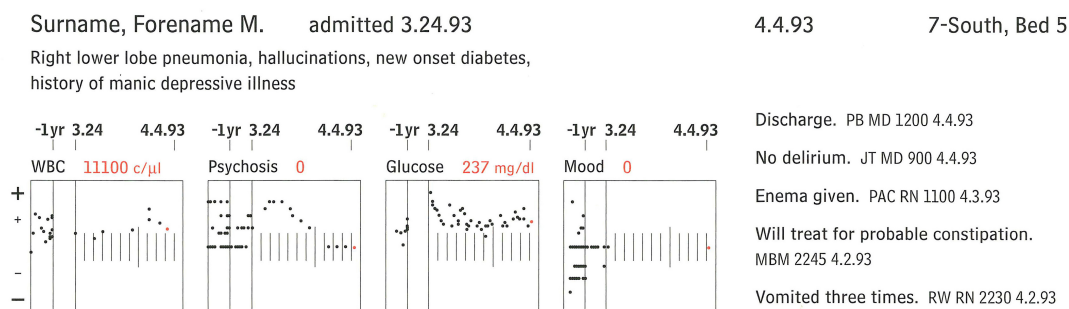


Figure 3.24: Part of the summary produced by the Graphical Summary of Patient Status [71]

3.2.2 VIE-VISU

VIE-VISU [46] 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 3.25). The motivation for this representation was to find a visualization supporting the immediate identification of stable patients, continuous improvement or worsening of the patient, as well as drastic changes in the patient's status.

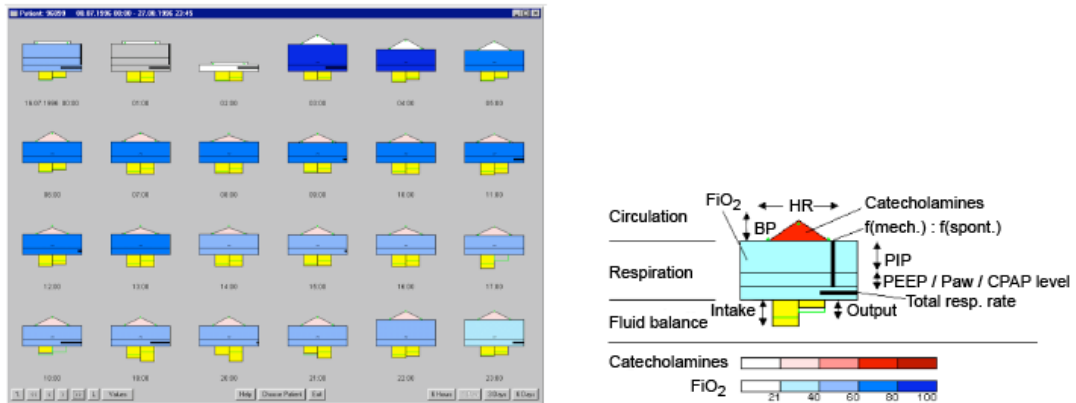


Figure 3.25: 24 hours display of VIE-VISU along with the legend of the glyph [46].

3.2.3 Interactive Parallel Bar Charts (IPBC)

Interactive Parallel Bar Charts have been designed by Chittaro et al. [17] 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. IPBCs represent such a time series in a bar chart format, where 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 lay out the time series of different hemodialysis sessions next to each other (see Figure 3.26).

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, Chittaro et al. 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. In order to solve a common problem of 3D representations, namely the occlusion of elements, Chittaro et al. provide a possibility to flatten occluding elements.

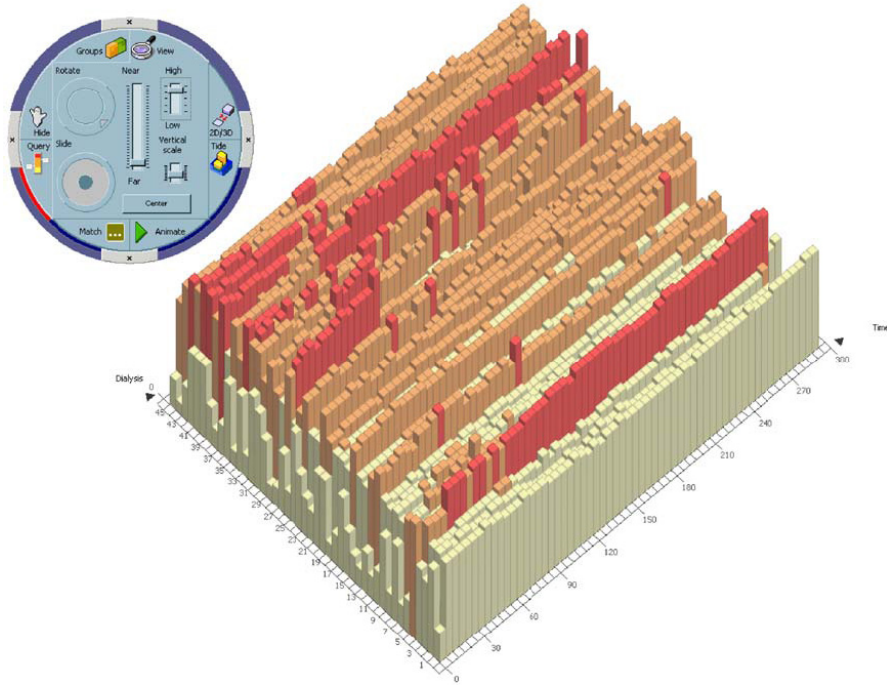


Figure 3.26: Interactive Parallel Bar Charts [17]: the X-axis (at the right) represents the time, the Y-axis (the height of a bar) represents the recorded values of one parameter during a hemodialysis session, and the Z-axis (at the left) shows the different hemodialysis sessions.

3.2.3.1 KNAVE and KNAVE-II

KNAVE [79] has been developed for the interpretation, summarization, visualization, explanation, and exploration of time-oriented clinical data. KNAVE-II [80] is an extension of KNAVE which supports the interpretation of time-oriented data as well as the visual exploration of these 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 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 are represented by line charts while qualitative data are represented by qualitative regions (see Figure 3.27). Moreover, interactions for synchronizing the displays of the panels according to time, adding statistics to the panels, zooming in and out, etc., are provided.

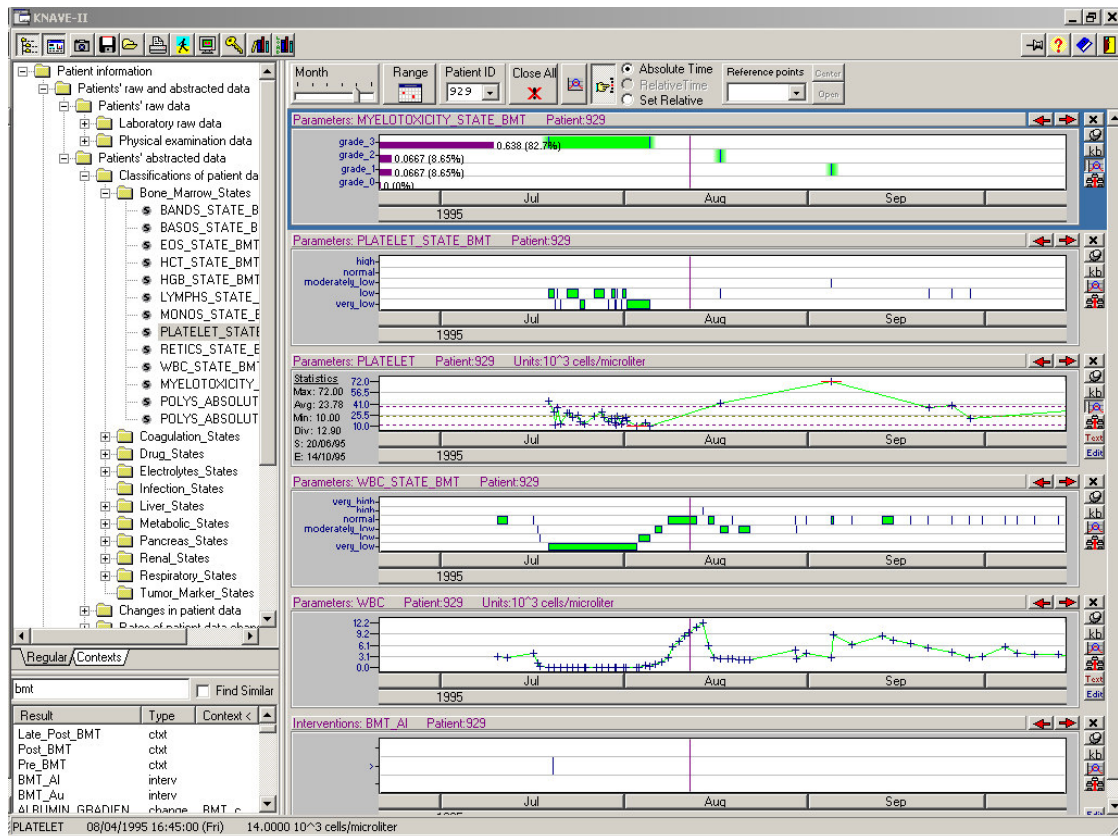


Figure 3.27: KNAVE-II [80]: different parameters of one patient can be explored with the help of multiple features, like different abstraction levels, absolute and relative time scales, different granularities of time, additional statistics, etc.

3.2.3.2 VISITORS

VISITORS [51] 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 Pan & Zoom, filtering of data, and Details on Demand facilitate the exploration process.

3.2.3.3 Gravi++

Gravi++ [43] 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. Psycholo-

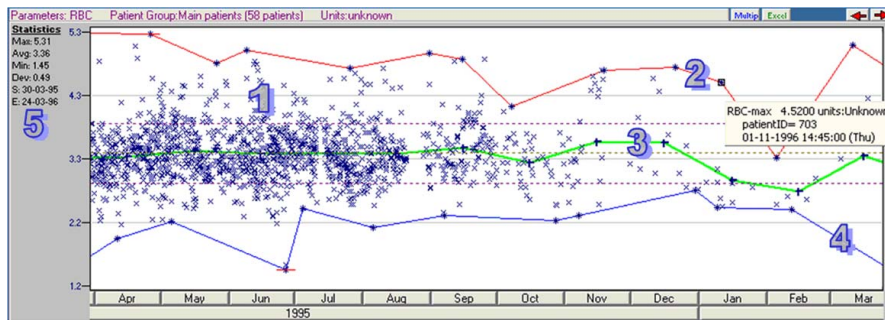


Figure 3.28: VISITORS [51]: visualization of a single raw parameter for a group of patients (58 patients) over time. The monthly minimum and the monthly maximum values of the group are indicated by the bottom line (blue) and the top line (red); the monthly mean value is represented by the green line in the middle. The three dotted lines indicate the mean value \pm standard deviation for the whole period

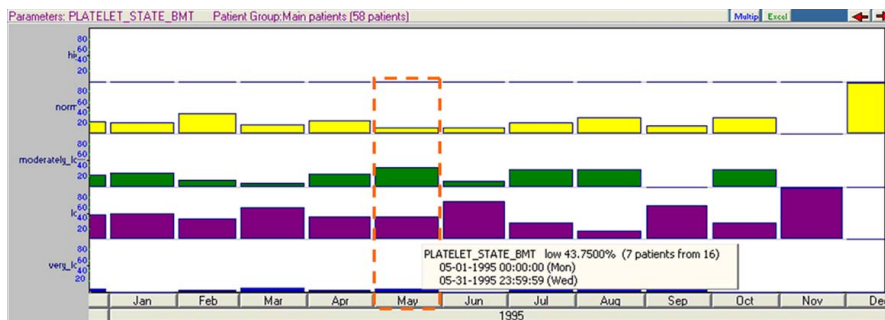


Figure 3.29: VISITORS [51]: visualization of one abstract parameter for a group of patients (58 patients) over time. They use a modified version of bar charts. The qualitative value is encoded by color. The heights of the bars segments indicate the percentage of patients of the selected group with respect to the qualitative parameter. The tooltip says '7 patients from 16' which means that there are 7 patients with the given qualitative value and only 16 patients in the group of 58 patients had the necessary raw data during May 1995.

gists and physicians need to derive time-oriented quantitative and qualitative information about the states of the patients to assess possible developments according to the past.

Gravi++ uses two kinds of icons to represent patients and questions. The questions' icons are arranged in a circle. A spring-based approach [23] is used to arrange patients' icons within this circle as follows: each patient is connected 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 rating of an answer (on a scale from '0' to '6'), the stronger the force that pulls the patient's icon towards the question icon. Thus, a patient's answers influence the position of the patient's icon on the screen, which leads to the formation of clusters of patients who have given similar answers (see Figure 3.31).

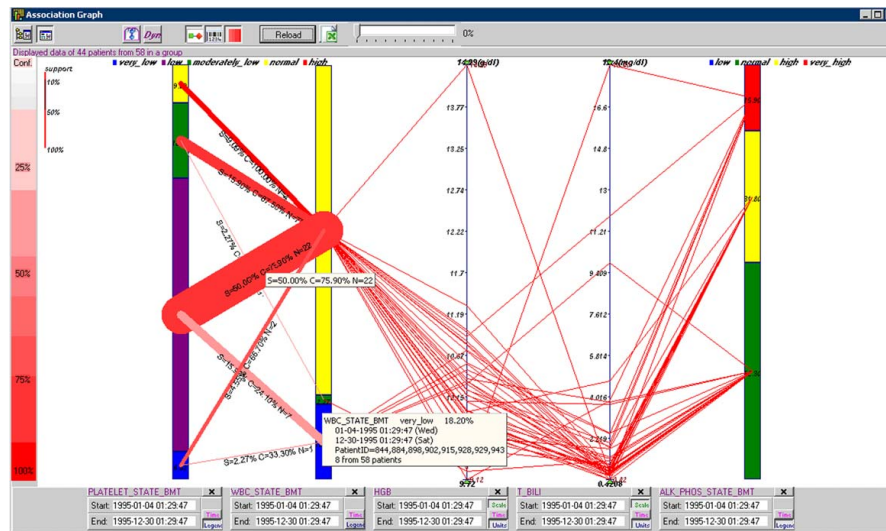


Figure 3.30: VISITORS [51]: visualization of multiple parameters (raw and abstract) for a group of patients (58 patients). Again, qualitative values are indicated by color and, additionally, in textual form (in the legend above the axis). The qualitative axes show the relative proportion the qualitative groups. For each parameter, the mean values of one patient (during a selected period of time) are connected by lines. The thickness of the lines between two qualitative axes indicates the number of patients with these two qualitative (mean) values.

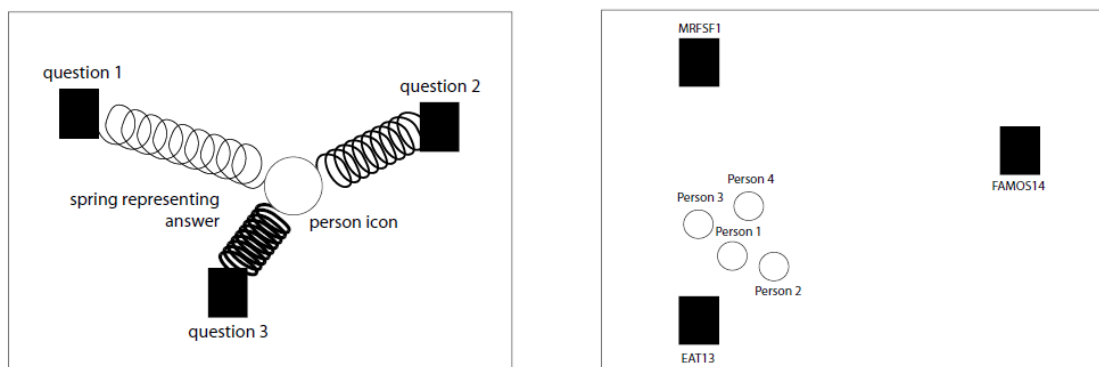


Figure 3.31: Spring based visualization: A patient is connected with each question by (invisible) springs. The higher the rating of an answer (on a scale from '0' to '6'), the stronger the force that pulls the patient's icon towards the question icon. This leads to the formation of clusters of patients who have given similar answers [43].

Moreover, the size of a patient'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' parameters over time, Gravi++ uses animation and traces.

Rings around the questions' icons along with Star Glyphs [15] 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 patient's therapy state. Red indicates a negative therapy development and green: indicates a positive therapy development (see Figure 3.32).

Interactive features, for instance changing the position or strength of a question or highlighting a question or a patient, facilitate the exploration of the data.

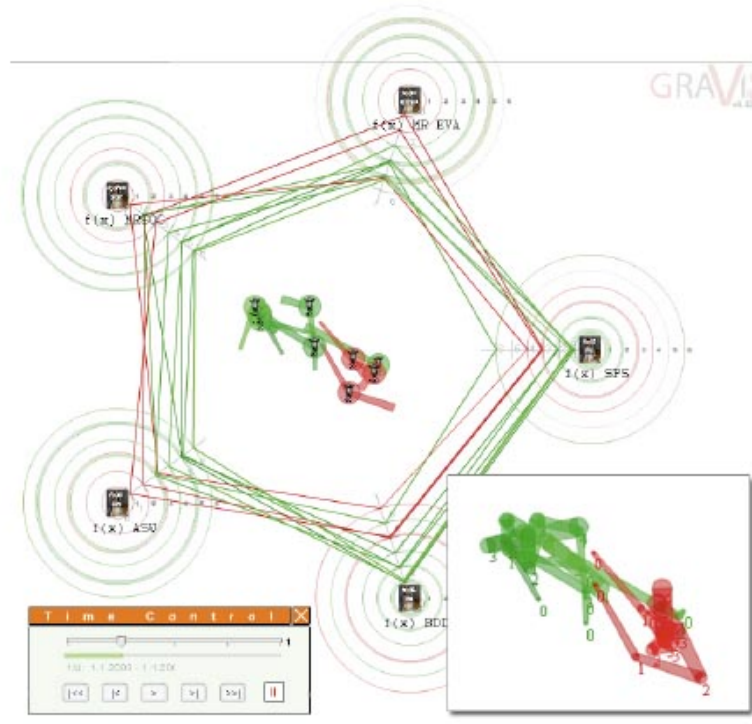


Figure 3.32: Gravi++ [43]: the course of these two patient groups (positive and negative therapy development) over time is visualized by animation and traces (as shown in detail on the lower right).

All these approaches of visualizing patient data have their specific benefits. However, a discussion on how well they are meet our requirements can be found at the end of this section (see Section 3.4). One crucial point is how well the visualization of patient data can be combined with the visualization of medical treatment plans.

3.3 Visualizing Medical Treatment Plans

The complex nature of medical treatment plans in terms of clinical practice guidelines demands for a plain and compact visualization of the underlying information. Dealing with logical se-

quences and hierarchical data as well as time-oriented data, the visualization of clinical treatment plans relates to several specific fields of InfoVis.

In recent years different frameworks to implement computerized guidelines have been developed, such as Asbru, EON, GLIF, GUIDE, Prodigy, and PROforma [4, 22, 64]. For each framework, various visualization methods have been developed in order to communicate the complex underlying information.

3.3.1 Clinical Algorithm Maps

Clinical Algorithm Maps are also called Flow Chart algorithms and are the most widely used representation of clinical guidelines or treatment plans (see Figure 3.33). 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 [29]:

‘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.’ [29, p.149]

This standard includes a small number of different symbols as well as some rules on how to use them. One extension to standard Flow Charts are annotations including further details: citations of 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 require 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.

3.3.2 Protégé

Protégé [31] is an extensible Java tool for the development of customized knowledge-based systems. The flexible 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, PROforma [86]). The graphical user interface illustrates the clinical algorithm in a way similar to Flow Charts by using different shapes for plans, decisions, actions, inquiries, and root tasks which are connected by arrows (see Figure 3.34).

3.3.3 Tallis Toolset

Being the successor of Arezzo[®] [47] - a software to generate, visualize, and execute PROforma [86] guidelines - the Tallis Toolset [86] is designed to create, publish and enact PROforma applications over the web. The PROforma task ontology is a set of reusable components: plans,

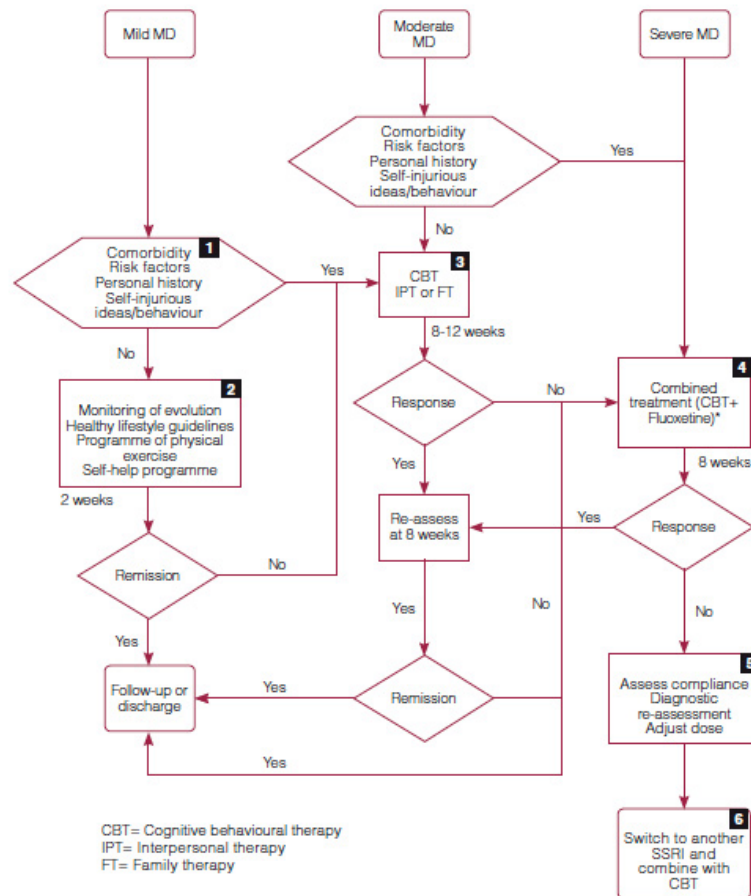


Figure 3.33: A Clinical Algorithm Map representing the logics of handling major depression in childhood and adolescence [20].

decisions, actions, and inquiries (see Figure 3.35). These tasks are composed to form networks (called ‘workflow’) and filled with details. All tasks have attributes, like goals, preconditions, and postconditions.

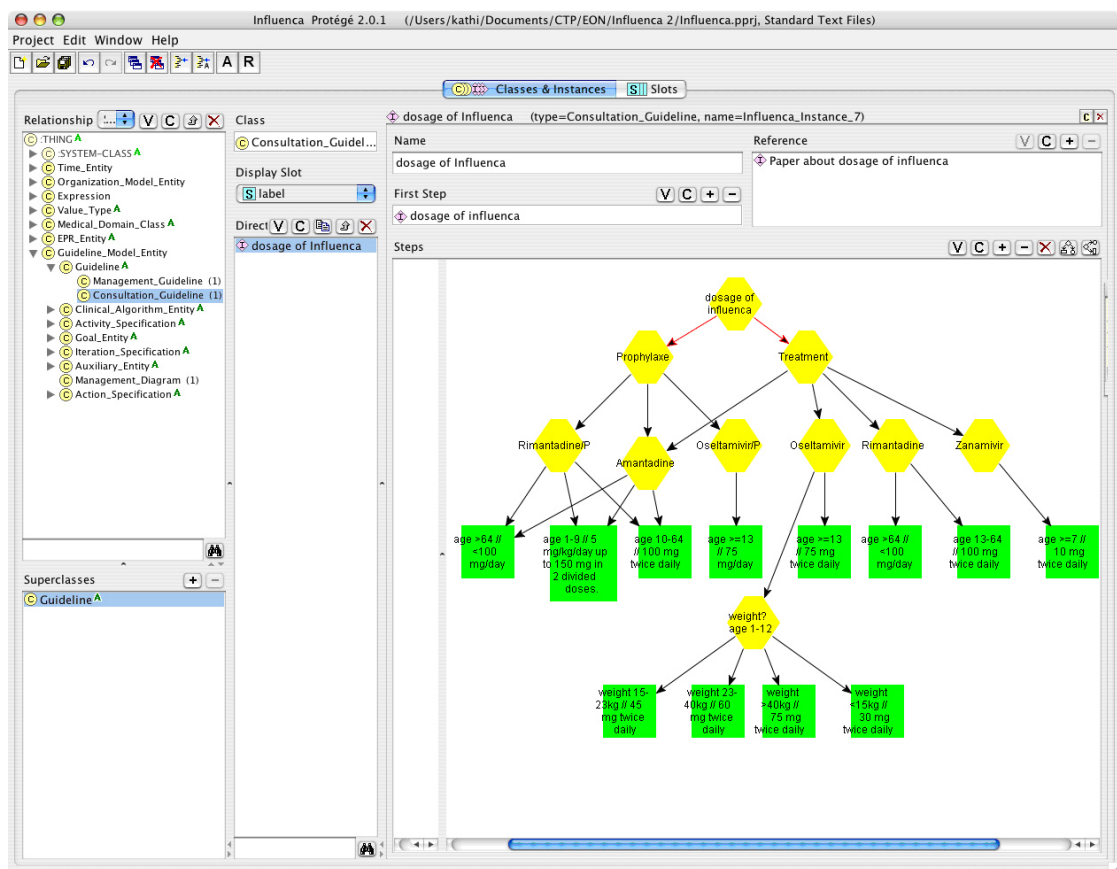


Figure 3.34: Protégé: Modeling a CPG about Influenza.

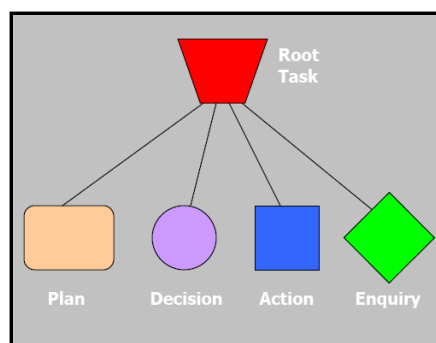


Figure 3.35: PROforma tasks: Tallis uses different shapes and colors to represent the different PROforma tasks [86].

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 3.36). 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.

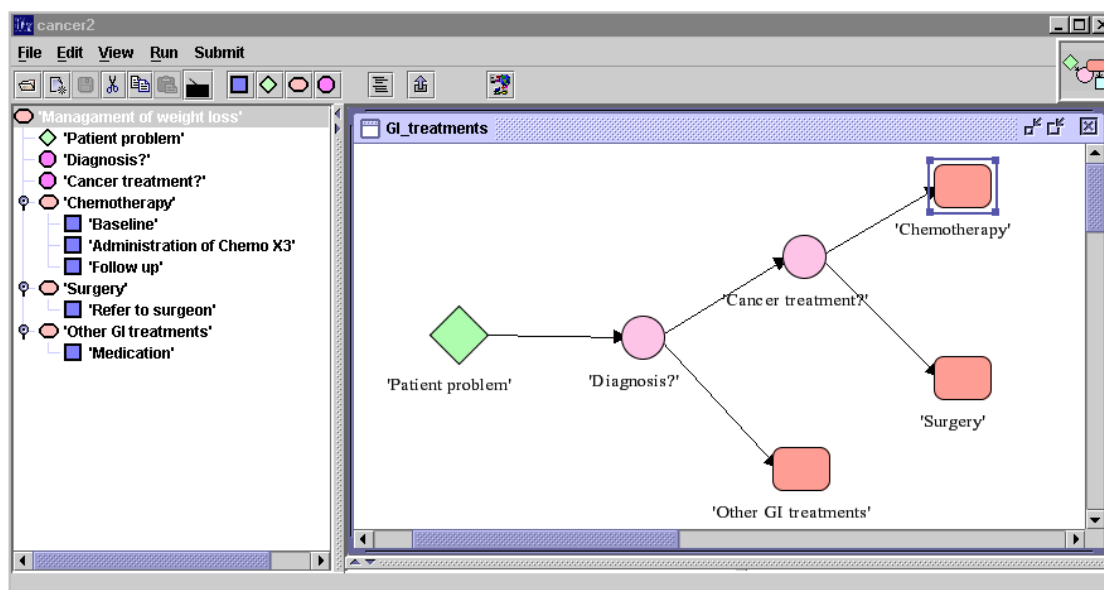


Figure 3.36: Tallis authoring environment: a simplified task network for cancer diagnosis and treatment [86].

3.3.4 GLARE

GLARE [88] (**G**uide**L**ine **A**cquisition, **R**epresentation, and **E**xecution) is a domain independent method. Similarly 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 provides primitives for drawing the guideline workflow similar to the representations of the Tallis Toolset and Protégé (see Figure 3.37).

3.3.5 GUIDE

GUIDE [72] 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 a modeled guideline into clinical workflow, using decision trees and influence diagrams to visualize complex relations, and using extended Petri Nets for the simulation of guideline implementation. Additionally, the GUIDE method allows for graphically authoring the workflow of guidelines. NEW-GUIDE [18] is the successor of GUIDE and is intended to facilitate the whole patient management, i.e., medical actions and organizational actions that can be managed by a workflow management system.

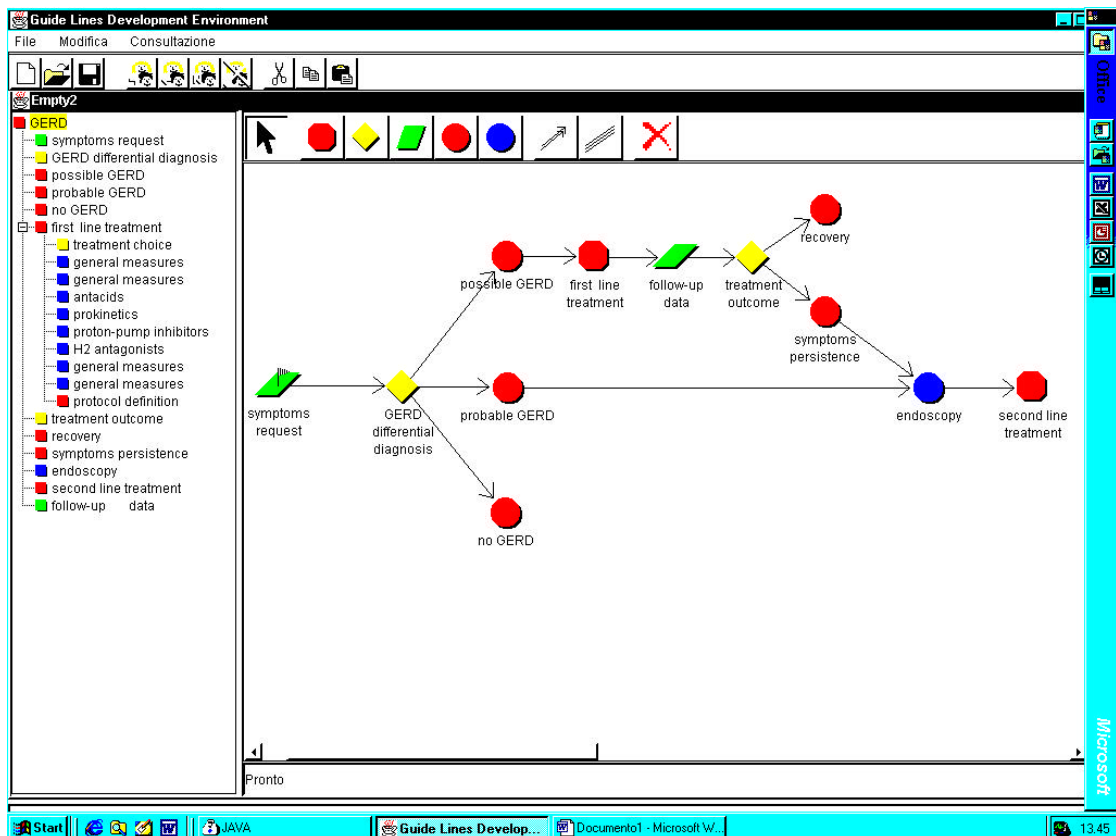


Figure 3.37: GLARE acquisition module graphical interface [88].

The NEW-GUIDE editor allows for graphically authoring the workflow of guidelines in a Flow Chart-like representation (see Figure 3.38).

3.3.6 VisiGuide (Part of DeGeL)

VisiGuide is part of the DeGeL (Digital electronic Guideline Library) project [82] which is a web-based architecture aimed at facilitating the transformation of a textual guideline into a formal model. The VisiGuide tool 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.

3.3.7 Guideline Overview Tool (GOT)

The main purpose of the Guideline Overview Tool [1] is to provide a compact overview of patient data in connection with applied treatment plans and to allow for the comparison of the progress of several patients at one view. GOT uses an extension of LifeLines [70] to visualize the various states of Asbru plans [58] as well as their execution order. Moreover, plans in the

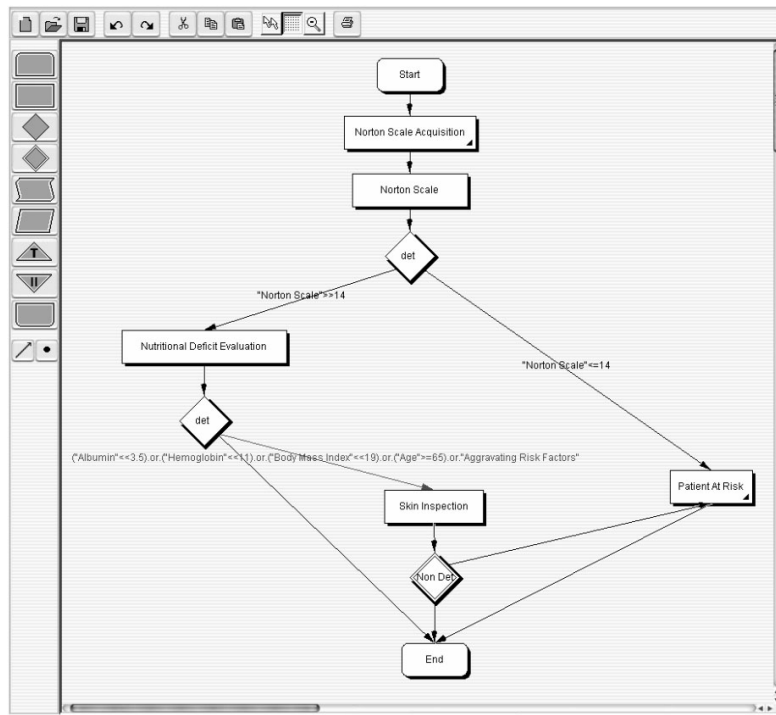


Figure 3.38: The NEW-GUIDE editor [19]: The guideline is modeled in a Flow Chart-like representation using different symbols.

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 is at the top of the interface; the largest pane shows treatment plans and patients' parameters and is located beneath the time scale. Additional data about the patients (variables, constants, and current values) are displayed at the right pane (see Figure 3.40).

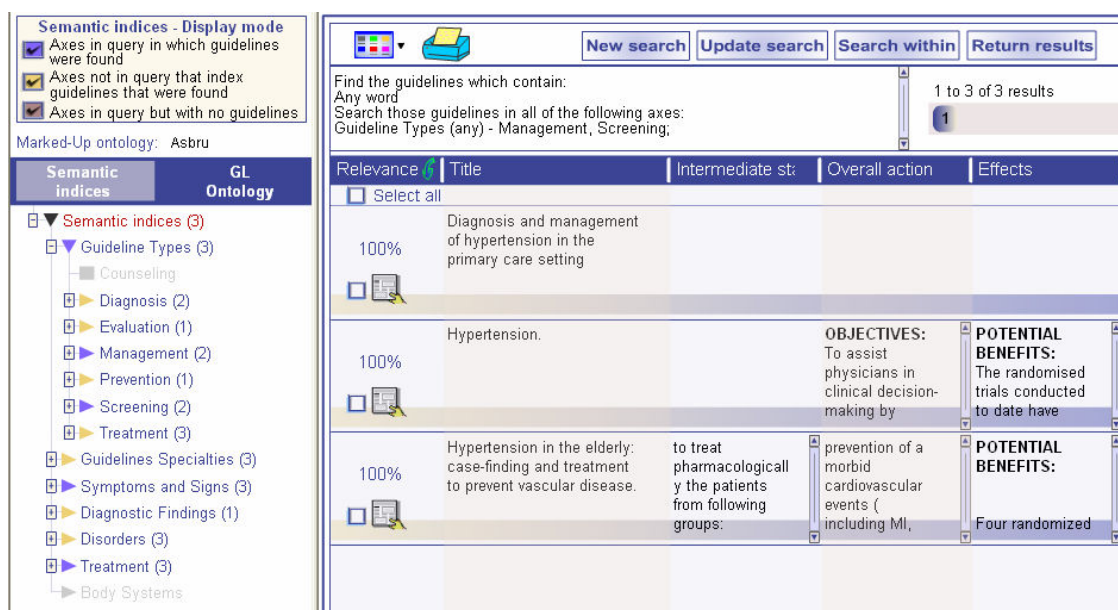


Figure 3.39: The VisiGuide tool is used to browse and present large amounts of guidelines as well as exploring a single guideline in detail [82].

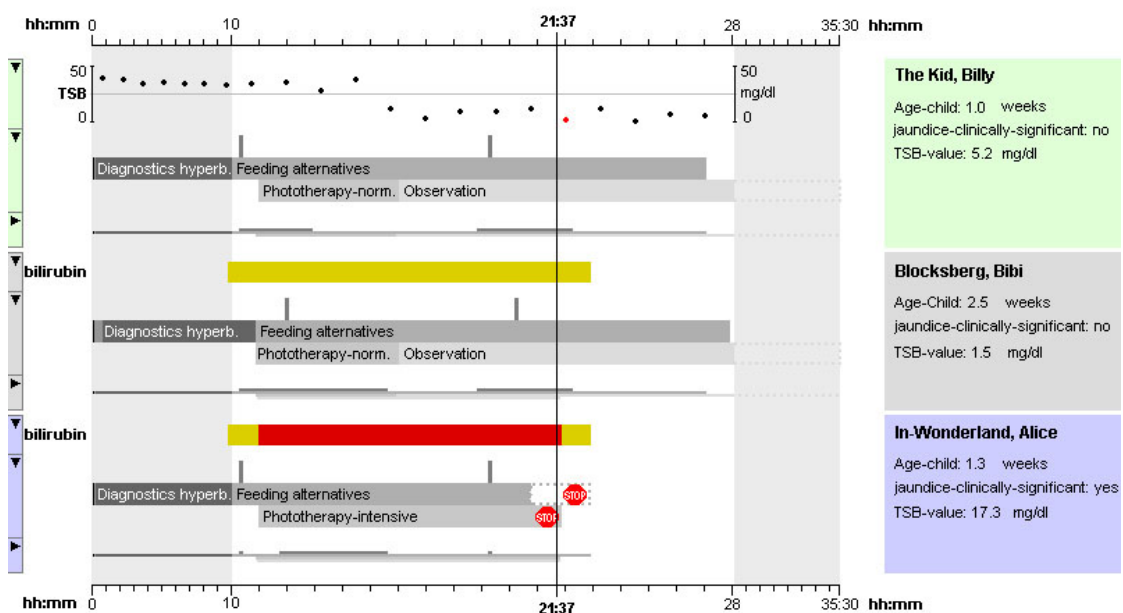


Figure 3.40: Guideline Overview Tool (GOT): visualization of the course of three patients treated with the same guideline for Hyperbilirubinemia [1].

3.3.8 AsbruView

AsbruView [52] is part of the Asgaard/Asbru project [81]. 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 3.41).

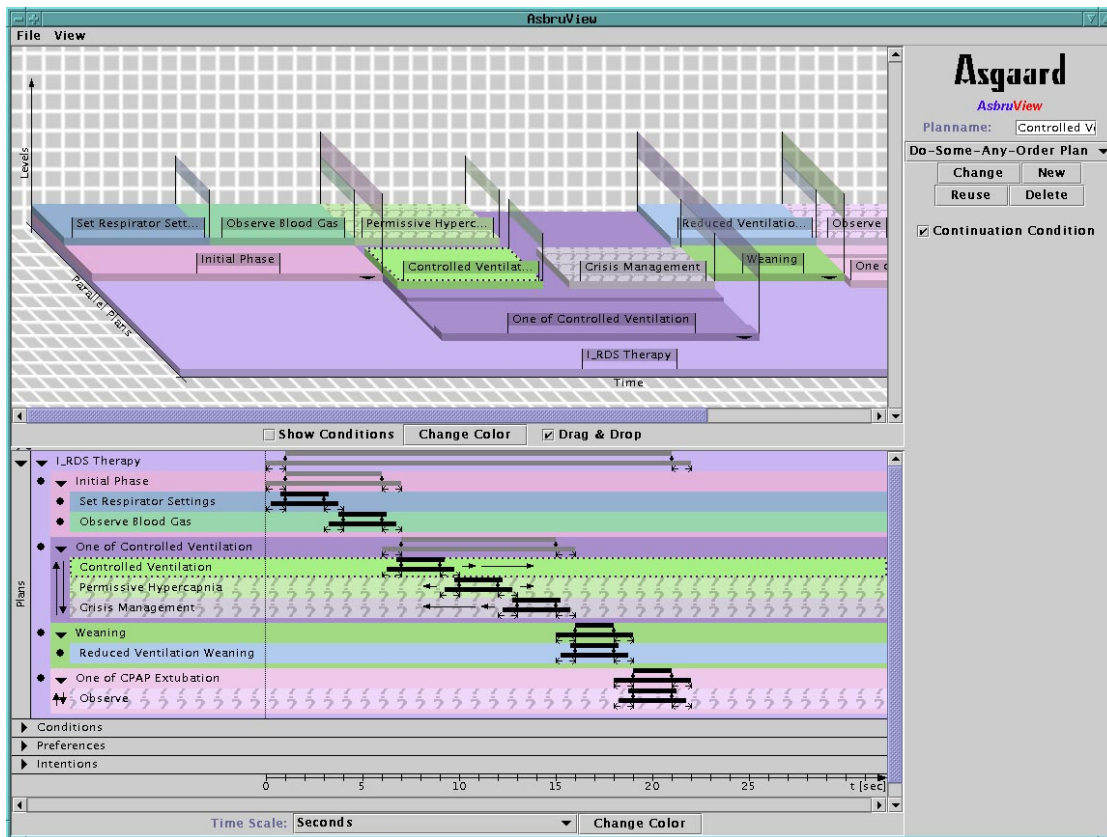


Figure 3.41: AsbruView: The Topological View (at the top) visualizes the composition of plans by means of traffic sign metaphors and the Temporal View (at the bottom) communicates complex temporal constraints [52].

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 of time. The central metaphor of the Topological View is a running rack representing an Asbru plan – the physician is considered to move 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 one after the other parallel to the time axis. Sub-plans of a parallel plan are put on the containing plan side by side 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 3.42). 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.

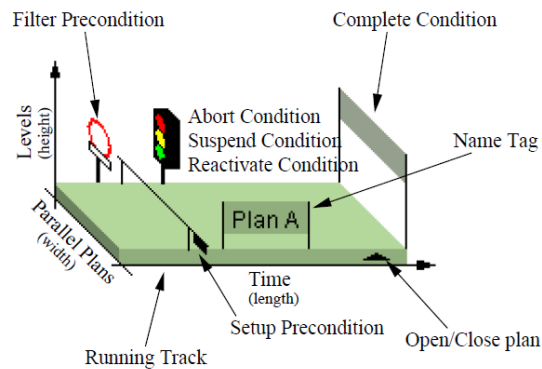


Figure 3.42: Anatomy of an AsbruView Plan in Topological View [52].

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 structure. The type of plans, i.e., the execution order of their sub-plans, is indicated by small symbols to the 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. These two diamonds support the second bar, which represents the minimum duration (MinDur) of the plan. If either the LSS or the EFS is 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 could not rest on 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 3.43).

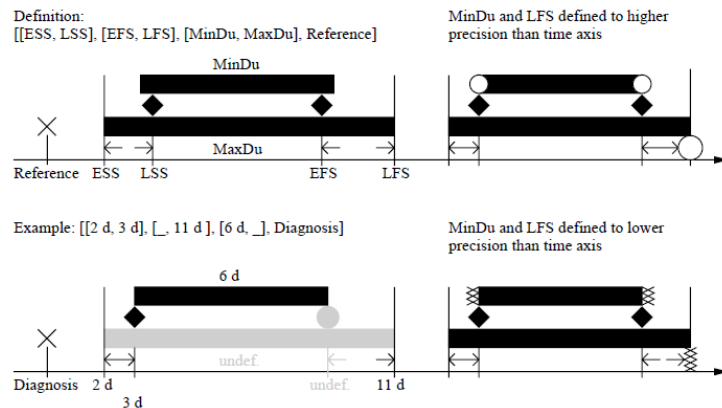


Figure 3.43: The Temporal View uses glyphs to communicate the exact temporal dimensions of the Asbru plans. The lower bar of the time-glyph indicates the maximum duration (MaxDur) while the upper bar indicates the minimum duration (MinDur) of the plan. The two diamonds on top of the MaxDur bar represent the latest starting shift (LSS) and the earliest finishing shift (EFS). The earliest starting shift (ESS), the latest starting shift (LSS), the earliest finishing shift (EFS), and the latest finishing shift (LFS) are indicated by vertical lines below the MaxDur bar [52].

3.3.9 CareVis

The CareVis interface [5] consists of three Views: in addition to the Logical View and the Temporal View, there is a Quick View panel (see Figure 3.44).

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 3.45). AsbruFlow is based on Flow Chart-like Clinical Algorithm Maps [40] 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, in addi-

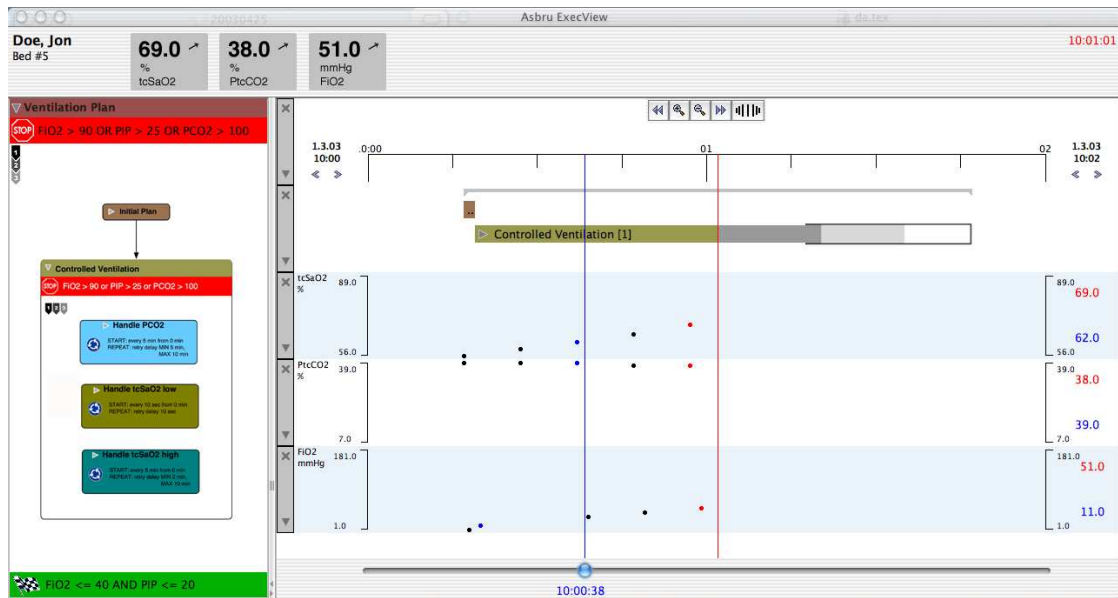


Figure 3.44: The CareVis application window showing the execution of a plan. The interface consists of three views: the Quick View panel (at the top), the Logical View (at the left), and the Temporal View (at the right) [5].

tion, the textual condition is given. 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 the possibility to collapse and enlarge sections, Focus & Context techniques, etc.

The Temporal View covers the representation of both time-oriented aspects of Asbru plans as well as parameters and variables over time. The visualized interval can extend 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 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 [70] providing additional features such as displaying the hierarchical structure 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 of future plans is called

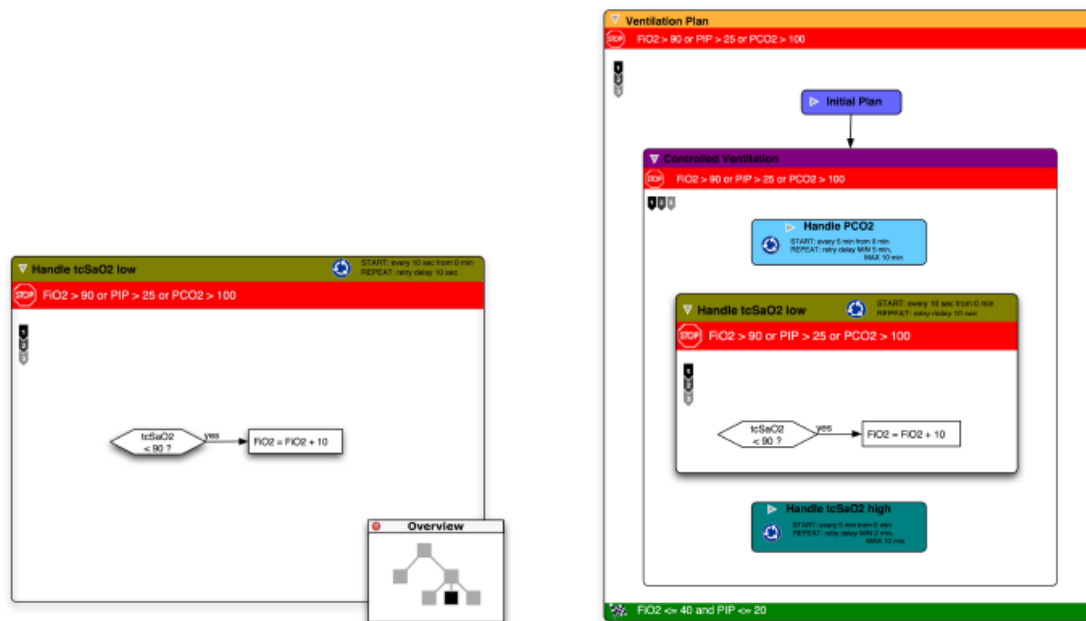


Figure 3.45: The Logical View (AsbruFlow) showing parts of the Asbru plan for artificial ventilation of newborn infants – Overview+Detail mode (left) vs. Fisheye view mode (right) [5].

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 3.46).

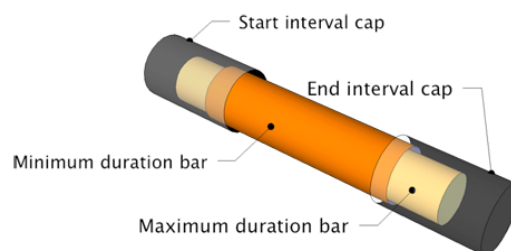


Figure 3.46: The PlanningLines glyph depicts the earliest starting shift, the latest starting shift (start and end of the start interval cap), the earliest finishing shift, the latest finishing shift (start and end of the end interval cap), the maximum duration, and the minimum duration of a treatment plan [7].

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 (see Figure 3.44 at the top).

3.3.10 Midgaard

Midgaard [10] is an interactive visualization prototype 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 color-coded qualitative regions which add additional knowledge to the presentation, for instance, if a particular value is critical or not (see Figure 3.47).



Figure 3.47: 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) [10].

They use the Semantic Zoom technique to brows through different visual representations and abstraction levels of the data from a broad overview to a fine structure of the data and several representations in-between (see Figure 3.48).

Bade et al. [10] introduce another visualization technique to overcome the problem of very limited space. When displaying high-frequency data, many data values would have to be depicted in 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 [92]. The space between the minimum and the maximum is 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 form a full data stream representation on very limited display space (see Figure 3.49).

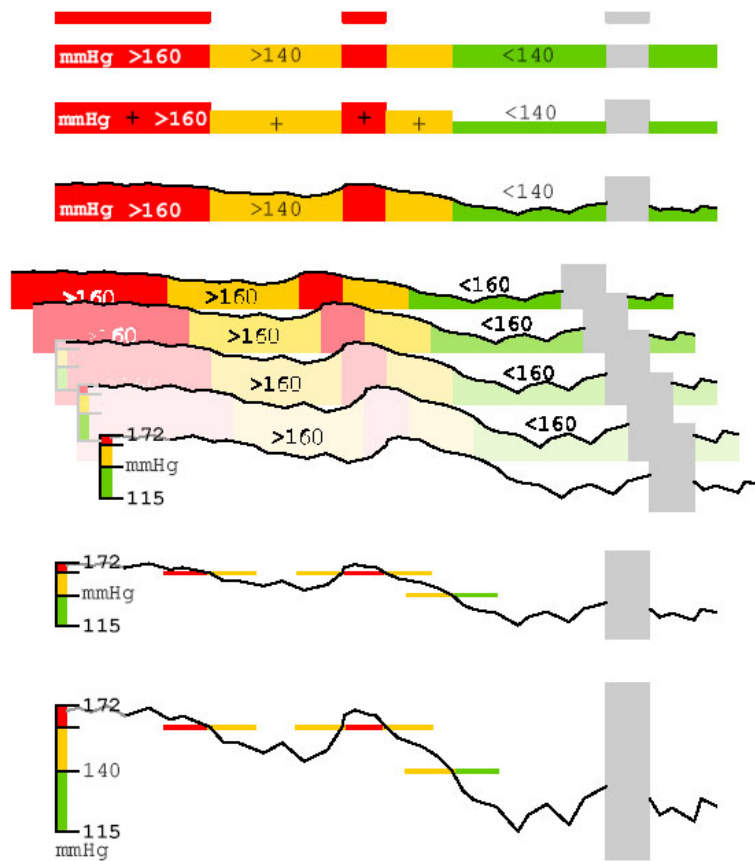


Figure 3.48: Semantic Zoom: browsing through different abstraction levels of the data (gray areas indicate missing values) [10].

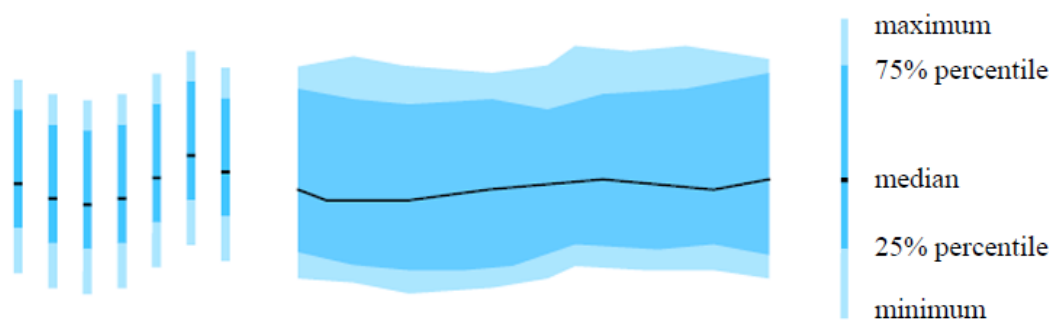


Figure 3.49: Redesign of the Tukey box-plot – inspired by Tufte [92]. Big amounts of these small box-plot-bars can be connected to form a full data stream representation of high-frequency data on very limited space [10].

Midgaard can also be used to visualize treatment plans. Figure 3.50 shows a treatment plan and its sub-plans. Plan characteristics such as preferences, conditions, intentions, and effects are given in tabular form.

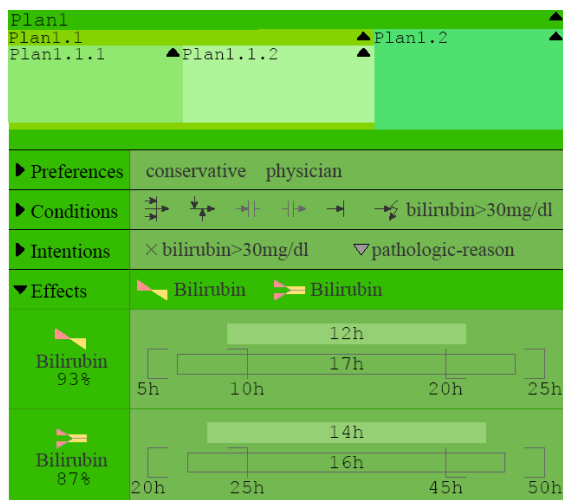


Figure 3.50: Visualization of a treatment plan as well as relevant plan characteristics (e.g., preferences, conditions, intentions, and effects) and the hierarchical structure of sub-plans. The estimated time of attainment of desired effects are visualized by special time-glyphs [9].

Moreover, the Midgaard prototype (see Figure 3.51) 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.

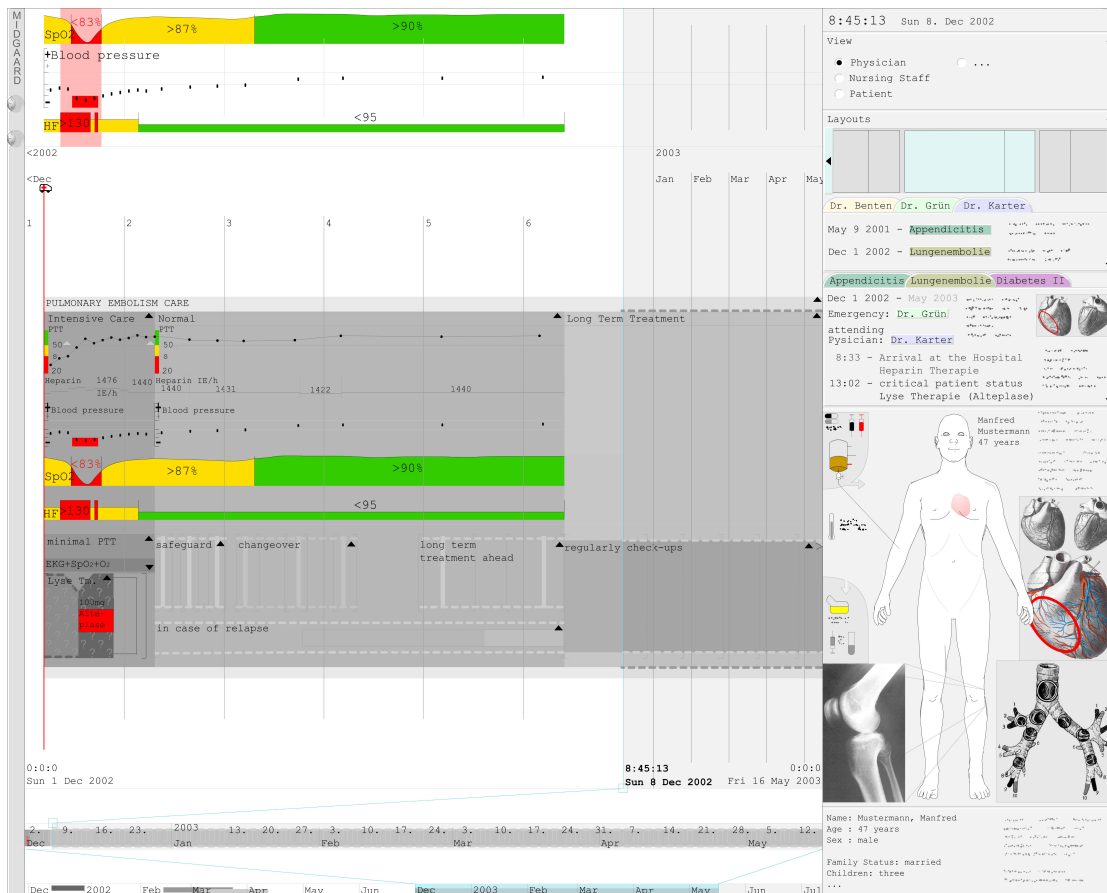


Figure 3.51: Screenshot of the Midgaard prototype [10].

3.4 Discussion and Shortcomings of Existing Approaches

In the previous sections we presented about 30 research approaches of visualizing hierarchical data, logical sequences, time series, patient data, and treatment plans. Next, we will outline how well the introduced visualization techniques meet the given list of requirements. See Table 3.1 for a compact overview.

Hierarchical data, logic sequences, and time series

Trees are a very intuitive way to represent hierarchical data, and so are Flow Charts, when it comes to logical sequences. In particular, physicians are familiar with Clinical Algorithm Maps which are quite similar to Flow Charts. On the other hand, Nassi-Shneiderman diagrams, Petri nets, and state transition diagrams are not well known among physicians and require some learning effort.

Now, the most important information we want to communicate (i.e., the effects of treatment actions on a patient's condition and the progress of patient parameters) is highly time-oriented.

PERT Charts give the temporal constraints more in a textual way than visually, which makes it hard to intuitively assess the temporal information. CircleView shows the qualitative change of a number of attributes but cannot be adapted to show treatment plans in combination with quantitative patient data. TimeSearcher uses conventional line charts while focusing on visual query tools.

CareVis already integrates trees and a Flow Chart-like representation to visualize hierarchical data and logical sequences. Moreover, it provides a Temporal View to represent complex temporal constraints of treatment plans by means of glyphs which are based on GANTT Charts and LifeLines.

The Line Graph Explorer uses color to reduce the Y-dimension of line charts to 1D-ribbons. In contrast to that, we decided for color-codings that take advantage of semantic information such as treatment plan goals and parameter progress since the start of the treatment. LiveRac uses background color to communicate one threshold value for the whole time window. However, we need to communicate more detailed information. In general, we use color in addition to the actual line chart to highlight and filter for specific events of interest as well as to establish a visual connection between clinical actions and line chart events.

Patient data

When taking a closer look at visualization prototypes that represent patient data, there are several interesting approaches. The Graphical Summary of Patient Status draws a comprehensive picture of a patient's condition over time, but it 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 user to explore them, while VISITOR is aimed at visualizing parameters of multiple patient records; however, both KNAVE and VISITOR neglect events that may influence the patient data. 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 requires 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 (namely occlusions) are still not solved satisfactorily. Gravi++ visualizes the change of the patient's condition over time by using animation and traces but it is not suited to giving a comprehensive overview of the patient's condition in one picture.

Treatment plans

There is a variety of approaches visualizing treatment plan information. Clinical Algorithm Maps, the Tallis Toolset, Protégé, GUIDE, and GLARE offer a Flow Chart-like representation of clinical algorithms suitable for visualizing the logics and execution sequence of plans, but not to visualizing either the hierarchy of treatment plans or complex plan characteristics. The purpose of VisiGuide is to browse and present a large amount of guidelines and to allow the user to skim through these guidelines and to explore a single guideline in detail. It displays

the guideline content and different characteristics, like conditions and intentions in textual form. However, there is no actual visualization of guideline plans. Then there is AsbruView, which succeeds in visualizing all relevant plan characteristics, but does not visualize patient data.

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, and 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 all the relevant characteristics of treatment plans but lacks a clear visualization of changes in the patient parameters due to treatment choices as well as a practicable assistance to explore the effects of applied clinical actions. Moreover, it does not support any inter-patient comparison.

In general, none of the mentioned approaches provides insights into effects of (multiple instances of) specific events (e.g., the repeated administration of a specific drug) on given parameter charts (e.g., the patient's condition), nor do they provide means to interactively investigate these effects.

	Hierarchical structure	Logical sequence	Time series (i.e., the patient's condition and its changes)	Temporal constraints of treatment plans	Detection of patterns	Comparing time series (i.e., multiple patients)	Characteristics of treatment plans
Trees and Treemaps	•						
Timeline Trees	•			•			
Flow Charts		•					
Nassi-Shneiderman diagrams		•					
Petri nets		•					
State transition diagrams		•					
GANTT Charts		•		•			
Time Lines and LifeLines		•		•			
PatternFinder				•	•	•	
PERT Charts		•		•			
CircleView			•		•	•	
TimeSearcher			•		•	•	
LiveRac			•			•	
Line Graph Explorer			•		•	•	
Graphical Summary of Patient Status			•				
VIE-VISU			•				
Interactive Parallel Bar Charts			•				
KNAVE and KNAVE-II			•				
VISITORS			•			•	
Gravi++			•		•	•	
Clinical Algorithm Maps		•					
Protégé, Tallis Toolset, GLARE, and GUIDE		•					
VisiGuide		•					•
Guideline Overview Tool	•	•	•	•			
AsbruView	•	•		•			•
CareVis	•	•	•	•			•
Midgaard	•	•	•	•			•

Table 3.1: Comparison of different visualization approaches as to how well they are suited to meet our list of requirements – outlined in Section 4 (•...the visualization approach meets the requirement).

Part II

CareCruiser

Requirements

In order to better understand the kinds of information that is required for our interactive InfoVis method, let us first take a look at a typical use case example of a treatment plan.

Use Case

An example application of a treatment plan would be the ventilation of premature new born children, i.e., neonates (see Figure 4.1). One part of this treatment plan aims at stabilizing the patient's transcutaneously assessed oxygen saturation ($tcSO_2$) and the patient's pressure of carbon dioxide (PCO_2). This is done by adjusting the supply of inspired oxygen (FiO_2) as well as adjusting the ventilation frequency (f). In particular, the intentions of the treatment plan are to stabilize the patient's $tcSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 $mmHg$.

$$\begin{aligned} 90\% \leq tcSO_2 \leq 92\% \\ \text{AND} \\ 40 \text{ mmHg} \leq PCO_2 \leq 60 \text{ mmHg} \end{aligned}$$

Moreover, the FiO_2 is to be kept below 40% and the peak inspiratory pressure (PIP) below 60 $cm H_2O$.

$$\begin{aligned} FiO_2 \leq 40\% \\ \text{AND} \\ PIP \leq 60 \text{ cm } H_2O \end{aligned}$$

This treatment plan can be decomposed into sub-plans, each of which may aim at achieving selected aspects of the overall intentions. Sub-plan 'Handle $tcSO_2$ ' aims at stabilizing the patient's $tcSO_2$ between 90% and 92% by adjusting the supply of FiO_2 (see Figure 4.2).

$$\begin{aligned} \text{IF } tcSO_2 < 90\% \text{ THEN } FiO_2 == FiO_2 + 10\% \\ \text{ELSE IF } tcSO_2 > 92\% \text{ THEN } FiO_2 == FiO_2 - 5\% \end{aligned}$$

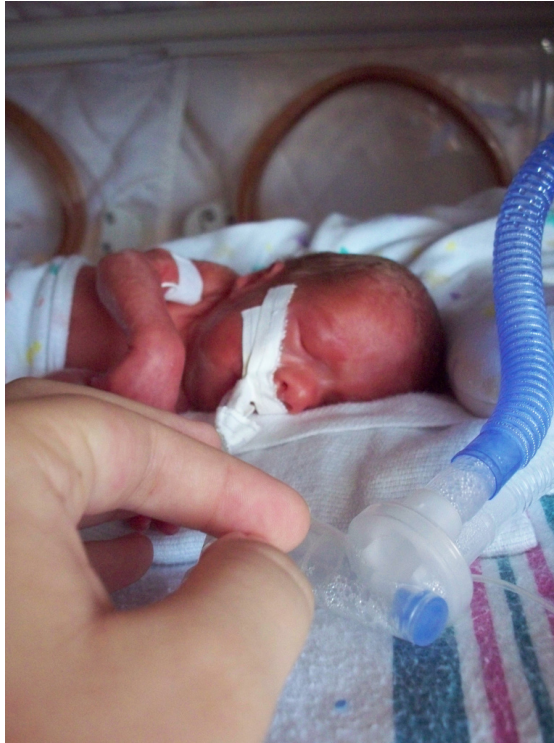


Figure 4.1: A premature infant requiring artificial ventilation at the neonatal intensive-care unit [98].

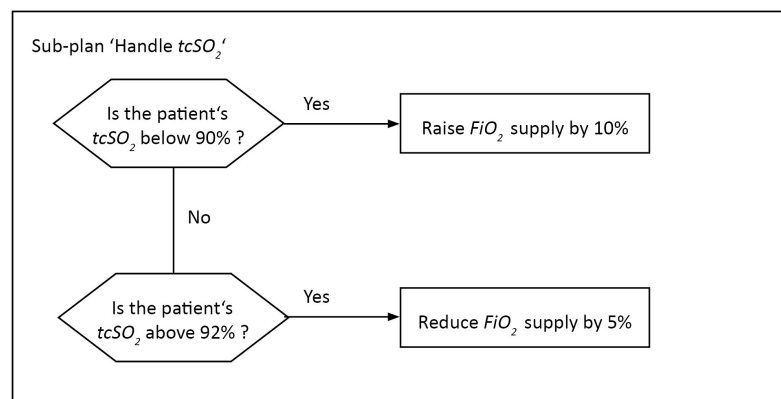


Figure 4.2: Sub-plan 'Handle $tcSO_2$ ' illustrated as a Clinical Algorithm Map.

Objectives

Our interactive InfoVis method aims to support the physician in monitoring the progress of one or more patients toward treatment plan intentions and in immediately spotting critical values at

the first sight. Moreover, the visualization prototype is to support the exploration of changes in the parameter curve possibly caused by clinical actions (i.e., effects of clinical actions) in order to optimize these treatment plans.

List of Requirements

A meaningful visualization of the changes in the condition of a patient – in connection with the execution of clinical treatment plans – has to include the communication of all important data and information. With respect to the characteristics of the above mentioned data and information, we have to cope with the following challenges:

1. Visualizing the **hierarchical structure** of treatment plans and sub-plans
2. Visualizing the **logical sequence** of treatment plan execution
3. Communicating the **temporal constraints of treatment plans**
4. Visualizing the effects of treatment plan execution and clinical actions on the patient's condition and providing means for the **detection of patterns**
5. Drawing a comprehensive picture of a patient's condition and its evolution over time (i.e., time series) and allowing for the **comparison of time series** (i.e., multiple patients)
6. Visualizing relevant **characteristics of treatment plans** (such as plan states, conditions, and intended effects)

In order to provide a method that supports all of these options, we decided to extend the possibilities of CareVis [5] by means of

1. enhancing the visibility of the patient's condition and its changes,
2. supporting the detection of possibly important patterns of reactions to clinical actions (such as the administration of a drug), and
3. providing the option of comparing different patients or different treatment plans.

CareVis [5] basically covers the nesting of treatment plans, the execution order, relevant plan characteristics, and temporal constraints. It 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 (compare Section 3.3.9).

Extending these possibilities of CareVis resulted in adding value to an already quite powerful system by making provisions to support every single aspect important in the application of computer-executable treatment plans. - However, it turned out to be challenging to find ways to smoothly integrate the visual encoding of the desired additional information into the user interface of CareVis without interfering with the existing visual encodings.

The following chapter contains a comprehensive description of our visualizations and interactive features and how we have integrated them into the CareVis prototype in order to meet the given requirements.

Design and Implementation

We have designed and implemented an interactive InfoVis prototype, called CareCruiser. We have extended the CareVis prototype with features to enhance the visibility of the patient's condition and its changes, to visually explore the effects that the applications of clinical treatment plans have on a patient's condition, features that support the identification of patterns of effects, and the option of comparing different patients. In accordance with the use case example given before (see Section 4), we use examples from the domain of artificial ventilation to demonstrate our approach.

5.1 User Interface Composition

In order to be able to properly design the user interface, we analyzed what kind of InfoVis method was best suited to meet the requirements of medical personnel. To make the system as versatile as possible, particular attention was paid to incorporating all kinds of data handling options into the system. For the sake of clarity, we will begin with an overview of these options before explaining in detail how they were implemented and what they are good for.

The user interface provides:

1. Display features
 - a) Charts of the patient's parameter values¹
 - b) A graphical representation of treatment plans:
 - i. Hierarchical nesting of treatment plans¹
 - ii. Logical structure of treatment plans¹
 - iii. Execution order of treatment plans¹
 - iv. Temporal constraints of treatment plans¹

¹features adopted from the CareVis prototype

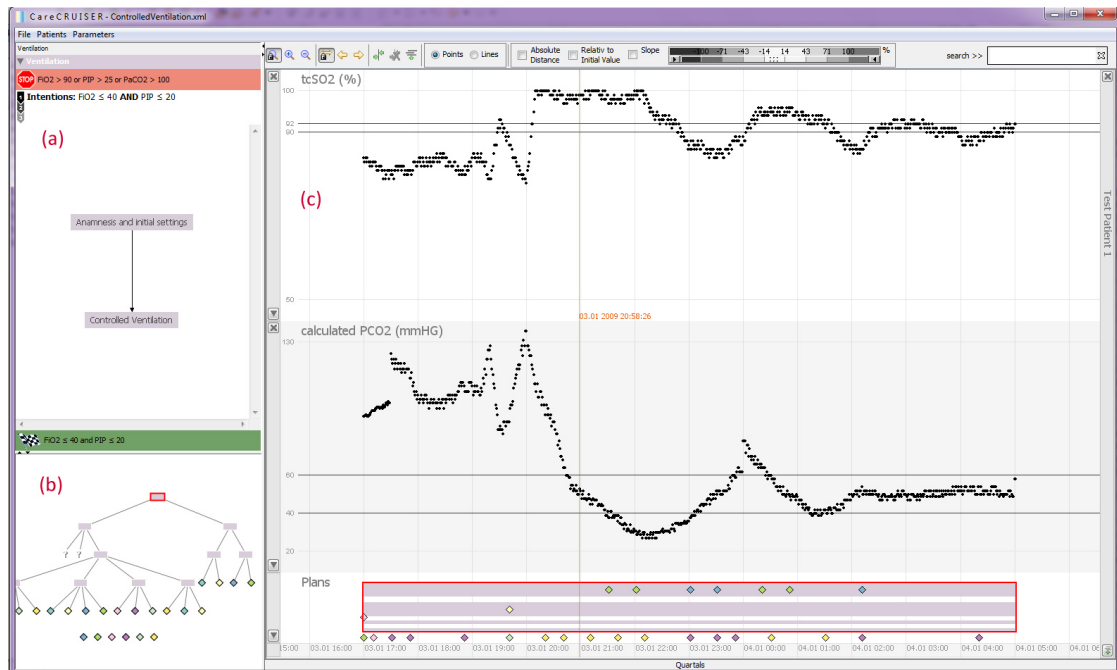


Figure 5.1: The CareCruiser user interface. The *Logical View* (a) communicates the logical structure of treatment plan execution by means of a flowchart-like representation [5]. The lower left part (b) displays a tree graph to visualize the hierarchical structure of treatment plans and sub-plans; the time-oriented view (c) focuses on the temporal-qualities of applied treatment plans, clinical actions, and patient parameters. We extended the time-oriented view with step-wise interactive means to explore the effects of applied treatment plans on the patient's condition.

v. Clinical actions and the time of application

2. Interactive means to compare selected treatment plans as well as selected actions in detail

- a) Color-coded distance information and slope
- b) Aligning plans or actions
 - i. *Distance to the intended value*
 - ii. *Progress from the initial value*
 - iii. *Slope of the parameter curve*
- c) Filtering the color-coded information
- d) A *Focus & Context Window*
- e) Further interactive features
 - i. Selection of patient parameters
 - ii. Highlighting of clinical actions
 - iii. Changing the parameter graph style

- iv. Details on Demand
 - v. Navigating the timeline¹
3. The option of comparing a certain number of patients

5.2 Visual Encoding

InfoVis uses visual attributes such as position, size, color, orientation, etc. to communicate complex ideas precisely [27]. Thus, the superior goal is to design an interactive InfoVis prototype that encodes all the required information and data in an easily understandable and consistent way. Visual artifacts 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.

1. Display features

The CareCruiser interface provides different coordinated views, glyphs, and charts to adequately visualize all relevant information about patient parameters and applied treatment plans.

(a) Charts of the patient's parameter values

We visualize the progress of a patient's parameter values (e.g., fever curves, oxygen saturation) as simple charts over time, which is a representation appreciated and well understood by clinical personnel (see Figure 5.2). 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 as physicians are used to assessing 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.

¹features adopted from the CareVis prototype

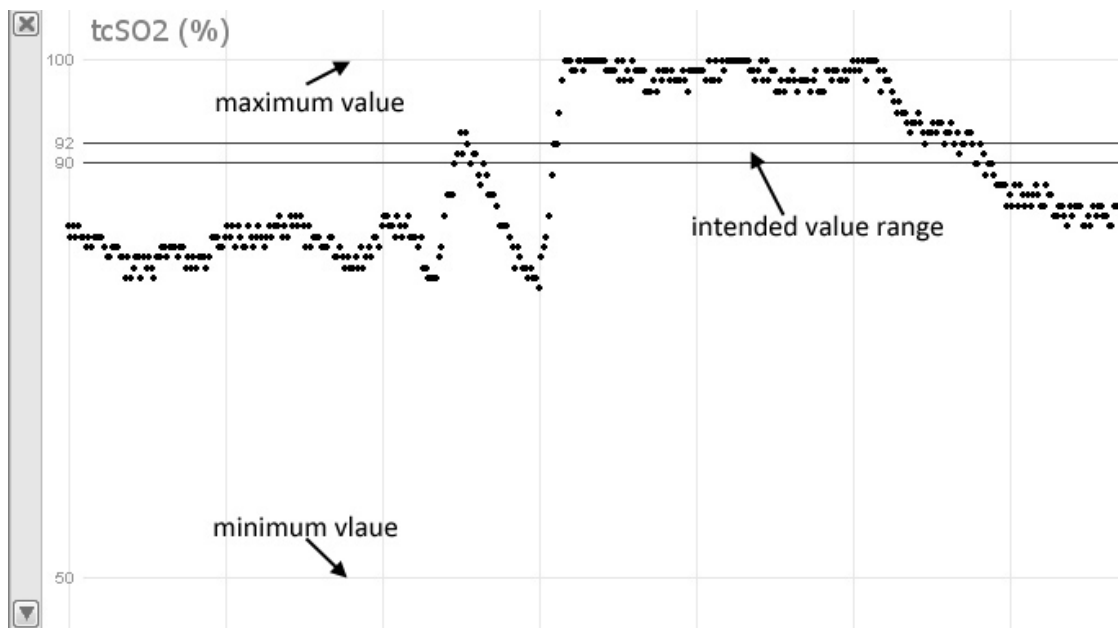


Figure 5.2: Chart of a patient's parameter value, the parameter being the transcutaneous oxygen saturation of the patient. The expected minimum and maximum values are indicated by the outer horizontal lines; the intended (i.e., desirable) value range is indicated by the inner horizontal lines.

(b) A graphical representation of treatment plans

When it comes to the visualization of treatment plans, there are several kinds of information that have to be communicated, such as the hierarchical structure of plans and sub-plans, the logics of a treatment plan (what plan is to be applied in what situation), starting- and ending-points of plans (usually there are some uncertainties for future plans), and what clinical actions are to be applied and when.

i. Hierarchical nesting of treatment plans As mentioned above, a typical treatment plan usually consists of a number of sub-plans, each of which may consist of further subplans in its turn. To represent such a hierarchical structure of a treatment plan and its sub-plans we provide a tree graph (see Figure 5.3) in the *Hierarchical View*, as tree graphs are the generally used visualization technique to communicate hierarchical structures. To facilitate the coordination between views (i.e., *Hierarchical View*, *Logical View*, and *Temporal View*), the selected plan elements are highlighted simultaneously in each view. We adopted the *Hierarchical View* from CareCruiser and extended it with specific glyphs for anamnesis and clinical actions, since clinical actions play a key role in visualizing the effects of treatment plans.

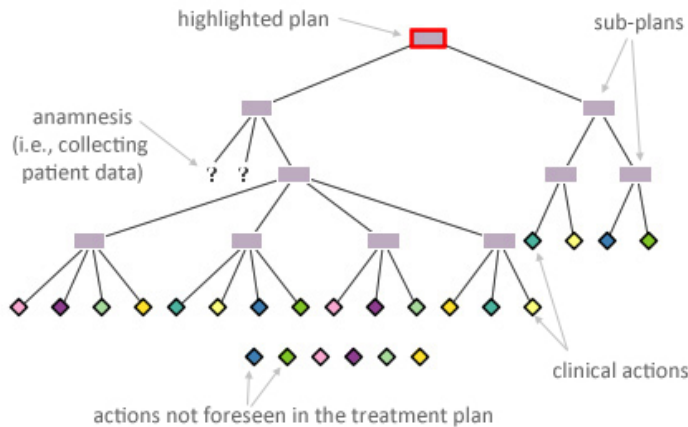


Figure 5.3: The *Hierarchical View* displays a tree graph to visualize the hierarchical structure of treatment plans and sub-plans.

ii. Logical structure of treatment plans The logical structure of a given treatment plan is visualized in the *Logical View* (see Figure 5.4). 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 a red stop-sign. At the top of the main window there is the execution sequence indicator and a textual description of the intentions specified for this treatment plan. The Complete Condition is represented by a checkered finish flag icon in the green bar at the bottom of the plan and also given in textual form (in Figure 5.4 no complete condition is specified for the given plan). Again, we adopted the *Logical View* from the CareVis prototype and just added diamond-symbols to clinical actions for the sake of consistency.

iii. Execution order of treatment plans With respect to the relevant temporal dimension, we arrange treatment plans along a timeline (the x-axis), this being the standard method of visualizing time. Other visual mappings are, for instance, animation, traces, and small multiples [93]. However, compared to static methods, animation leads to many participant errors; traces are less accurate than a small multiple display [75] and small multiples 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 communicates the starting- and ending-points of the plan as well as its duration on the horizontal timeline.

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. A plan might be the artificial ventilation of newborn children; in this case a sub-plan could be the stabilization of the child's oxygen saturation. In the case of sequential plans and parallel plans, the execution order of their sub-plans is strictly predefined.

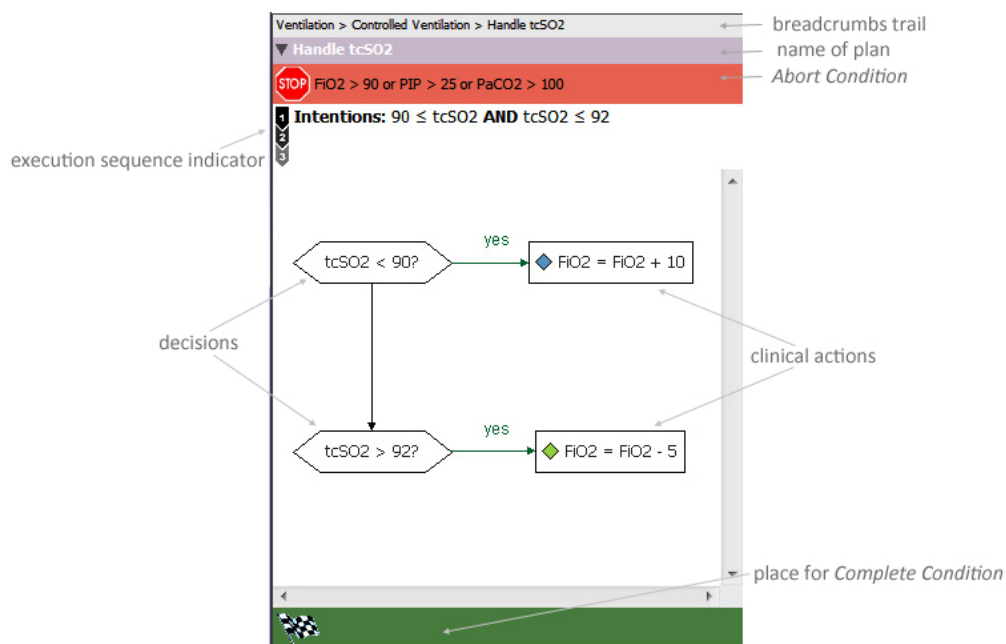


Figure 5.4: The *Logical View* communicates the logical structure of treatment plan execution by means of a Flow Chart-like representation [5].

- **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 5.5).

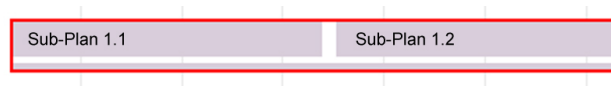


Figure 5.5: Sequential plan: the sub-plans are laid out on a timeline (x-axis) one after the other which means that the plans are carried out in sequence.

In addition, the execution sequence is indicated by different symbols in the upper left corner of the *Logical View* (see Figure 5.4). There are four possible symbols indicating the execution order (see Figure 5.6)



Figure 5.6: Execution sequence indicator: four different symbols indicating four different execution sequences [5].

- **Parallel Plans**

All sub-plans of a parallel plan are 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 their parallel execution (see Figure 5.7).

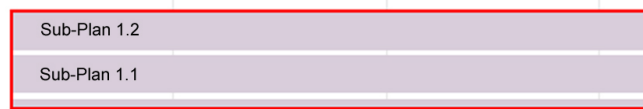


Figure 5.7: Parallel plan: The sub-plans are stacked vertically which is to indicate that the plans are carried out in parallel.

- **Cyclical Plans.** A cyclical plan contains a sub-plan that can be repeated. The time interval for which the plan is repeated, the starting point, the ending point, and the duration of the particular plan 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 5.8).



Figure 5.8: 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.

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 in 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.

iv. Temporal constraints of treatment plans The starting- and ending-points of plans is given implicitly by the outlines 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 above the parent-plan, which does not violate the time-axis (see Figure 5.5 and Figure 5.7).

Other representations of hierarchical data – for instance, trees and treemaps (see Section 3.1.1.1) – cannot be used in this context because of their inability to handle time. Timeline Trees (see Section 3.1.1.3) are a sophisticated method to visualize hierarchical data with respect to time, but they do add some complexity to navigation and fail in giving a complete overview of all hierarchy levels in one picture.

To integrate Asbru’s complex time annotations of *earliest starting shift*, *latest starting shift*, *earliest finishing shift*, and *latest finishing shift* as well as the maximal and minimal duration of a plan, we have integrated the CareVis PlanningLine glyph (see Section 3.3.9) into our visualization (see Figure 5.9).

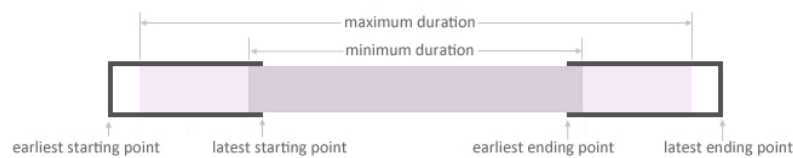


Figure 5.9: PlanningLine glyph [5]: the caps (the bracket-like forms at the far left and far right of the glyph) represent the earliest and latest starting shift and finishing shift; the maximum duration is indicated by the light gray area and the minimum duration is indicated by the dark gray area.

v. Clinical actions and the time of execution For assessing the effects of clinical actions the user needs to know when, how often, and in which context these actions have been taken. To visually represent these actions we chose diamonds for two reasons (see Figure 5.10):

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).

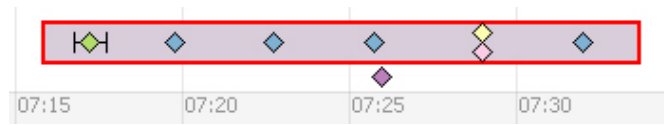


Figure 5.10: Clinical actions (e.g., the administration of a drug) are indicated by diamonds. The peak of the diamond indicates the exact point in time when the action is applied while the body of the diamond ensures the visibility of the action. In case an action is carried out over a time span, the temporal bounds are indicated by whiskers. Clinical actions that are applied to the patient but are not part of the treatment plan are laid out below the plan body.

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 starting- and ending-point by the width of the diamond would lead to line-thin diamonds for actions that take place at one single point in time, which would decrease the visibility of these actions. Similar to the representation of parallel plans, actions which are executed in parallel are aligned above each other. Based on the qualitative color palettes proposed by Harrower and Brewer [41], we assign a color 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 the context of executing this treatment plan. Moreover, the color of clinical actions facilitates the identification of one and the same action in different views. Since the number of well-distinguishable colors is limited, we additionally support the identification of clinical actions by tooltips and the possibility to highlight all applications of a selected clinical action. However, the number of different colors is sufficient for our test datasets.

Clinical actions that are applied to the patient but are not part of the treatment plan are laid out below the plan body.

2. Interactive means to compare selected treatment plans as well as selected actions in detail

In order to determine success or failure of an applied treatment plan, we take into account the intentions specified for this plan (intentions represent goals and intended effects of the treatment plan). To visually encode how well an applied treatment plan succeeds in reaching the specified intentions and to make the effects of applied clinical actions stand out more clearly, we use color-coding to highlight distance information and slope of the patient's parameter development. We opted 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. As already mentioned, physicians prefer to have these raw values to give them a first idea of the patient's condition. Apart from highlighting color-coded information, the step-wise investigation of effects of applied treatment plans and clinical actions is supported by the following interactive means:

- aligning plans or actions below each other to ease comparison (similar to the method used for aligning clinical events in LifeLines2 [95]),
- filtering for interesting events of parameter values' development, and
- selecting regions of interest by means of a *Focus & Context Window*.

The following subsections contain a detailed outline of these means.

(a) Color-coded distance information and slope

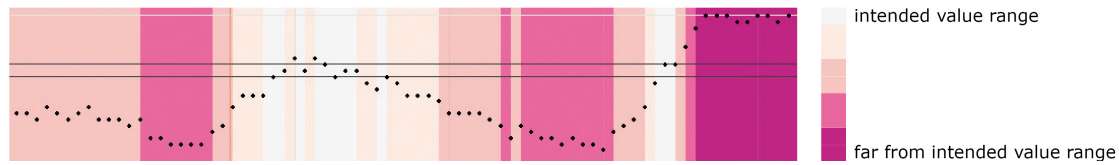
We have made it possible for the user to visually enhance three interesting aspects of a patient's parameter development, namely

- the *distance* of the parameter value to the *intended value*,
- the *progress* of the parameter value from the *initial value*, when the treatment started, and
- the rising and/or falling (i.e., the *slope*) of the parameter value.

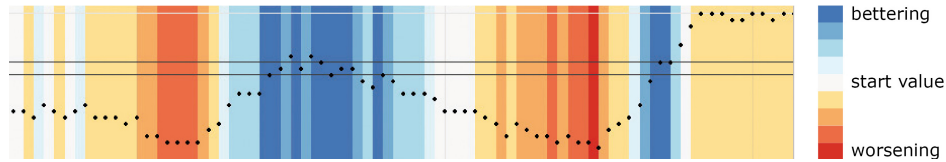
Each of these modes is suited to bring to light very different qualities of the data (i.e., the patient parameters in combination with the applied treatment plans). We use color to visually encode this information, because color is a strong visual attribute which gives an immediate impression of the varying extent of success or failure of applied treatment plans. We are aware of the problem that colors may have different meanings depending on one's cultural background; the color mappings we have chosen, however, are likely to be considered intuitive by most in the western world. In addition, we decided to have discrete color scales to make qualitative changes in the parameter development stand out more clearly.

i. Distance to intended value A patient's parameter value is affected by the applied treatment plan. The distance of this value to its 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 [41] from dark magenta to light magenta (see Figure 5.11(a)). 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.

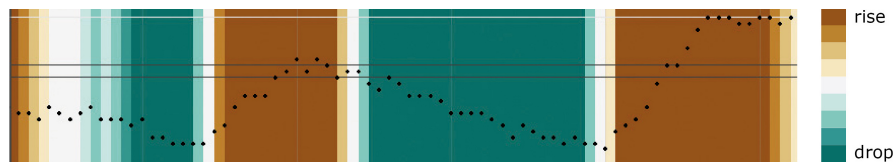
Figure 6.2 shows one treatment plan that has been applied to three different patients; the three instances are aligned vertically. The charts and treatment plans are colored according to the patient parameter values' *distance to the intended value*. In the figure, ranges far off the intended value are selected with the help of a range slider (see Paragraph (c)). Here, the visualization emphasizes the bad $tcSO_2$ values of the first patient. In contrast, the minor changes in the $tcSO_2$ curve of the third patient are close to the intended value range and do not draw off the user's attention. In this way, attention can be focused on critical cases only.



(a) *Distance to the 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 first sight.



(b) *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.



(c) *Slope*: Highlighting the *slope* of a parameter value (turquoise: drop, brown: rise). This mode helps to identify the immediate effects of applied clinical actions. In order to disregard minor oscillations of the parameter curve, we take the mean value of seven data points to compute the *slope*.

Figure 5.11: Comparing different modes of highlighting the effects of applied treatment plans. Each of them helps to reveal different aspects.

ii. Progress from initial value The main and evident criterion for the appraisal of treatment plans is how well the plan succeeds in achieving the intended effect. To visually encode this information, we consider the relative *progress from the initial value* (e.g., the oxygen saturation of a patient's blood 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 toward 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 effect.

This abstraction method (namely, calculating the progress of the parameter values on the basis of 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.

How important it is to define the success of a treatment with respect to the initial parameter value is illustrated in 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 [41] from blue over white to red. We map the success values from 0% (initial value) to 100% (intended value) with colors ranging 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 here, too, more saturated colors represent more extreme values (see Figure 5.11(b)).

Both the *distance to the intended value* and the *progress from the initial value* represent the patient's condition and the success of a plan with respect to one specific parameter.

iii. Slope of the parameter curve Visually enhancing the *slope* of the parameter chart, i.e., showing if the parameter values rise or drop, offers 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 toward or away from the intended value) in the parameter chart; we encode the falling or the rising of a *slope*. This is why we use a diverging color scheme [41] that is as neutral as possible: we represent negative slopes with turquoise highlighting and positive slopes with brown highlighting (see Figure 5.11(c)).

(b) Aligning plans or actions

To facilitate the comparison between the effects one treatment plan has had on two different patients or the effects of two alternative treatment plans, we provide a tool to align selected treatment plans together with their parameter curves below each other (see Figure 6.2). 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 (see Figure 5.12). 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 every single application of this action and subsequently comparing all of 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, thus making it easy to compare the effects the applications of this action have had on the patients' parameters.

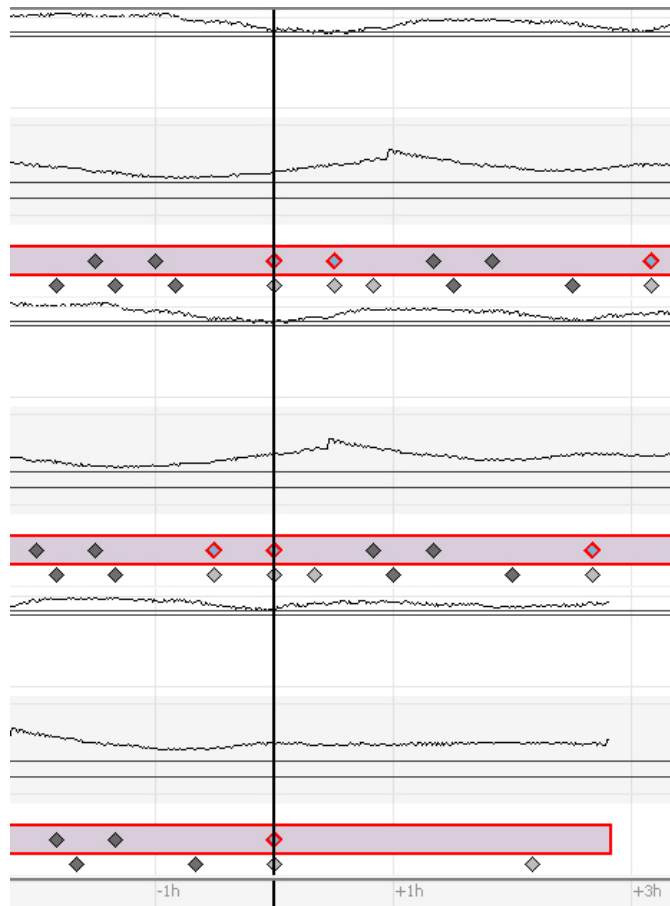


Figure 5.12: Aligning three instances of applying one clinical action to one patient. This helps to find out whether this action keeps having the same effects on given patient parameters.

(c) Filtering the color-coded information

The user may be interested only in specific curve events: for instance, critical values or extreme drops of the curve. Therefore, we provide a range slider filter to look for such events; the filter makes them stand out immediately (see Figure 5.13).

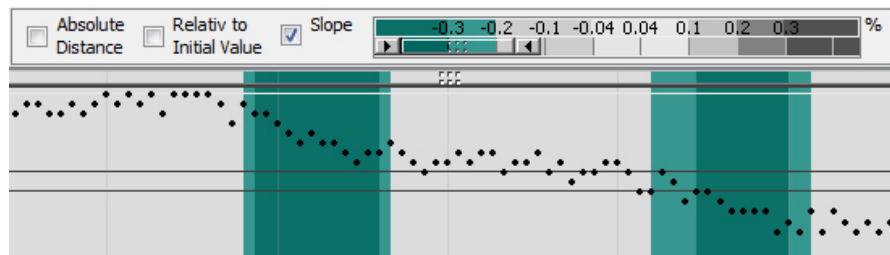


Figure 5.13: Selecting only extreme drops of the parameter chart with the range slider makes the coloring of lesser drops and rises of the chart disappear.

(d) A Focus & Context Window

After color-highlighting and filtering curve figure information and aligning plans or actions of interest, the user may investigate a specific region (e.g., the hour after the various applications of one clinical action). To support this task we have developed a *Focus & Context Window* which grays out the color-information outside its borders. The width of the window is varied by dragging its side borders, 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 & Context Window* are indicated by slightly darker regions, which maintains some context information.

(e) Further interactive features

Apart from the interactive means outlined in the previous section, we provide some minor interactive means to have a smoother workflow, customized perspectives, and easier access to diverse kinds of information.

i. Selection of patient parameters Besides the parameters specified within the treatment plan intentions, additional parameters of interest cannot be generalized; they depend almost entirely on the specific condition of the patient. Therefore, we provide means to add and remove patient parameters to/from the visualization. Thus the user is able to create customized perspectives of the visualization.

ii. Highlighting of clinical actions A clinical action of interest can be selected in the *Hierarchical View*. Subsequently, all applications of this action are highlighted within the applied treatment plans in the *Temporal View*. This relieves 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 mouse-over labels.

iii. Changing the parameter graph style The user can choose if the parameter values are represented (1) by single points which show the actual time of measurement, (2) by lines which

connect the points in time of parameter measurement and provide a more continuous representation – thus emphasizing trends, or (3) by a combination of both.

iv. 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 InfoVis prototype. 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 labels are an effective way to provide detailed information on single elements of interest while not interfering with the rest of the visual representation (i.e., without requiring a change of view) [84].

v. 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; also, one can 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 state of the patient, but also serves as a visual assistance and reminder of applying adequate treatment plans and clinical actions on time.

3. The option of comparing a certain number of patients

The CareCruiser prototype allows for the investigation of two or more patients simultaneously. The number of patients is limited to 3-4 patients with expanded patient facets on one screen (see Figure 5.14) or about 15 patients with collapsed patient facets (see Figure 5.15). The available display height is equally distributed to the patient facets. When moving the mouse over the right border of the patient facet, a pop-up window provides basic information about the patient such as name, age, gender, etc. (see Figure 5.14). However, when the patient facets reach a defined minimum height, the view becomes scrollable. To this end, collapsed patient facets allow for investigating a bigger number of patients at the same time (see Figure 5.15). These facets can be expanded, collapsed or closed individually or collectively.

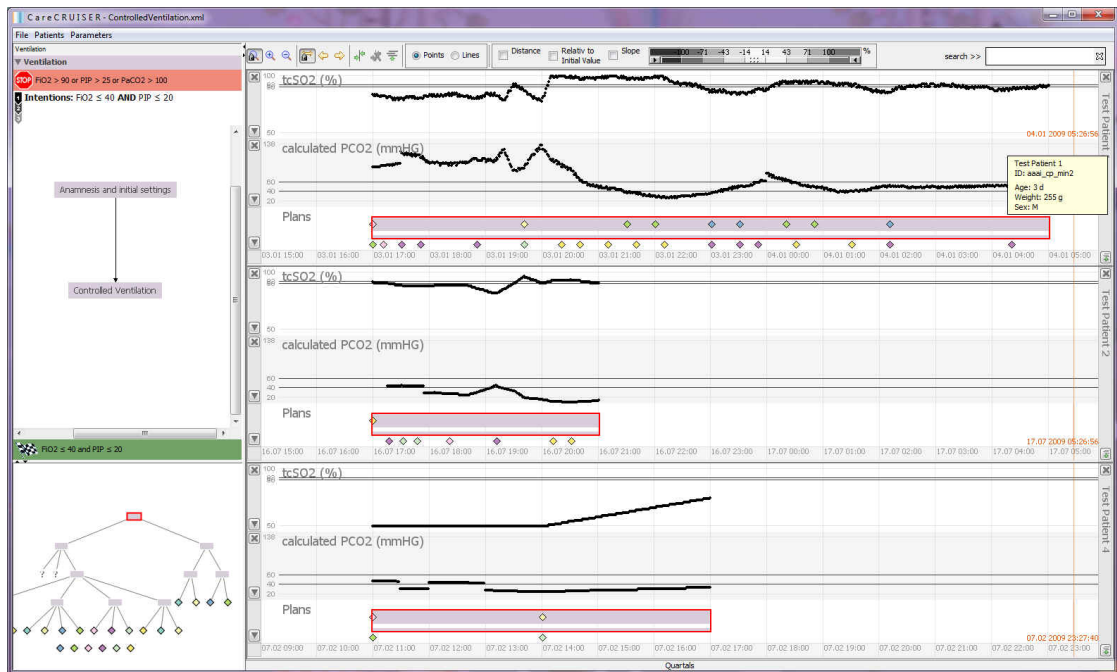


Figure 5.14: The same treatment plan has been applied to three different patients (aligned vertically to ease comparison).

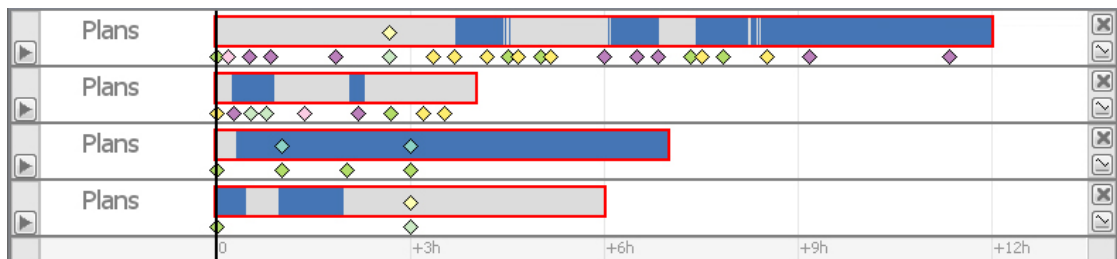
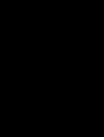


Figure 5.15: Collapsed patient facets. We have aligned a given treatment plan that has been applied to four different patients. When highlighting the parameter values' *progress from the initial value* and filtering for parameter values within the intended range (blue), one can easily see that the treatment plan succeeded in stabilizing the values of the first and third patients while it failed regarding the second and fourth patients.

In this section we introduced the different features of our interactive InfoVis prototype. Let us now illustrate these features with three clinical examples in the next section.



Comparing Effects and Discovering Temporal Patterns

There is more to visualization than meets the eye. It is common knowledge that visualizing allows you to intuitively grasp complex phenomena and developments at first sight without having to resort to laboriously dealing with a large amount of numbers and putting them into context. Visualizing, however, has more to offer: if properly employed, it has the potential to yield unexpected results that cannot be derived from algebraic methods. One of the more spectacular examples of this is Mitchell Feigenbaum's discovery of universality in deterministic chaos [32]. Feigenbaum had to rely on a computer to plot the graph that consisted of an enormous amount of dots representing the results of as many calculations. It was only with the help of this visualization that he could see and investigate the recurrence of patterns inherent in all such chaotic systems (see Figure 6.1).

With this example we just want to point out the usefulness and superiority of visualization over other methods of data mining in certain circumstances. This includes the fact that the visualization of data offers the most immediate feedback as to changing data.

Both CareCruiser and CareVis put these advantages of visualization to good use; CareCruiser being the patient-oriented complement of the more treatment plan-oriented CareVis. Actually, CareCruiser goes beyond CareVis by adding a number of dynamical extensions whose main object is the probing into the reaction of the patient to any kind of treatment that has been applied. Appropriate consideration has been given to the time factor, since this is an often neglected dimension that needs special consideration in the complex treatment plan and patient environment.

The following three examples illustrate how CareCruiser can be used in different clinical situations: (1) comparing the conditions of different patients, (2) providing a comprehensive overview over a patient's recuperation and how well the applied treatment plan succeeds, and (3) detecting temporal patterns in reactions to clinical actions.

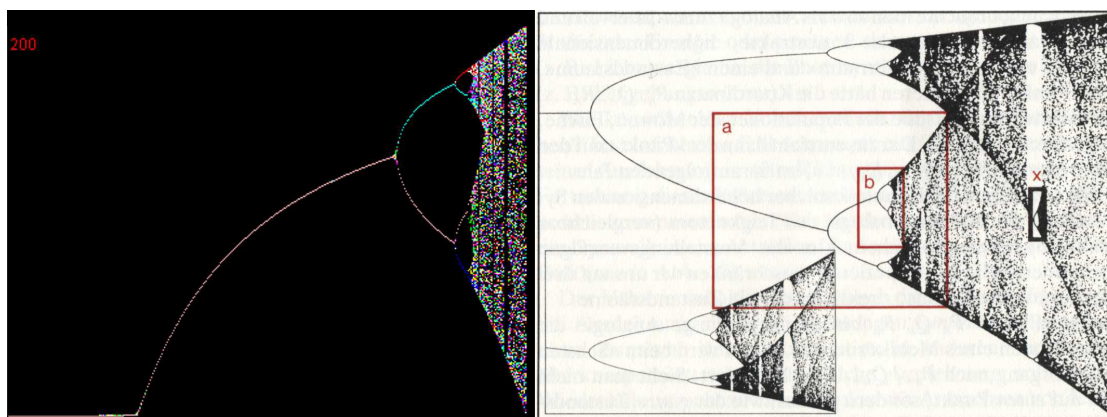


Figure 6.1: The bifurcation diagram. ‘A picture says more than a thousand calculations.’ Thanks to the fact that, with the help of his computer, Mitchell Feigenbaum succeeded in visualizing the results of hundreds of thousands of calculations, he could ‘see’ the conditions of deterministic chaos, which enabled him to discover a new constant of nature that was named after him. The left picture gives you an idea of what Feigenbaum might have seen on the monitor of his computer in 1975; the enlargement on the right side shows his discovery, namely that the graph exhibits the feature of self-similarity. Feigenbaum then found out that the ratio of the distances between these elements of self-similarity – from left to right – is a constant [24].

6.1 Comparing the Conditions of Different Patients

It is a standard situation in clinical institutions to treat patients having the same condition according to the same treatment plan. However, most of the time a treatment plan still has room for improvement; moreover, there usually exist several alternative plans to handle a given clinical situation. Visualizing the effects of one treatment plan upon a number of patients that suffer from the same illness yields fairly reliable results of how well this plan functions. In our example the differences in the patients’ conditions show up clearly. Since the differences in the health state of patients are emphasized by distance coloring, the critical values stand out (see Figure 6.2). Thus, using CareCruiser facilitates the monitoring, overviewing, and comparing of a number of patients.

Under the assumption that either the same hospital uses an alternative treatment plan as well, or that two different hospitals have alternative treatment plans and use CareCruiser to compare their results, the relative assets of a treatment plan as compared to its alternative will become evident, making CareCruiser instrumental to the eventual improvement of treatment plans and patient care.

6.2 The Success of the Applied Treatment Plan

This example has to do mainly with color and its quality to make immediately visible certain qualitative situations aimed at drawing a comprehensive picture about how well a treatment

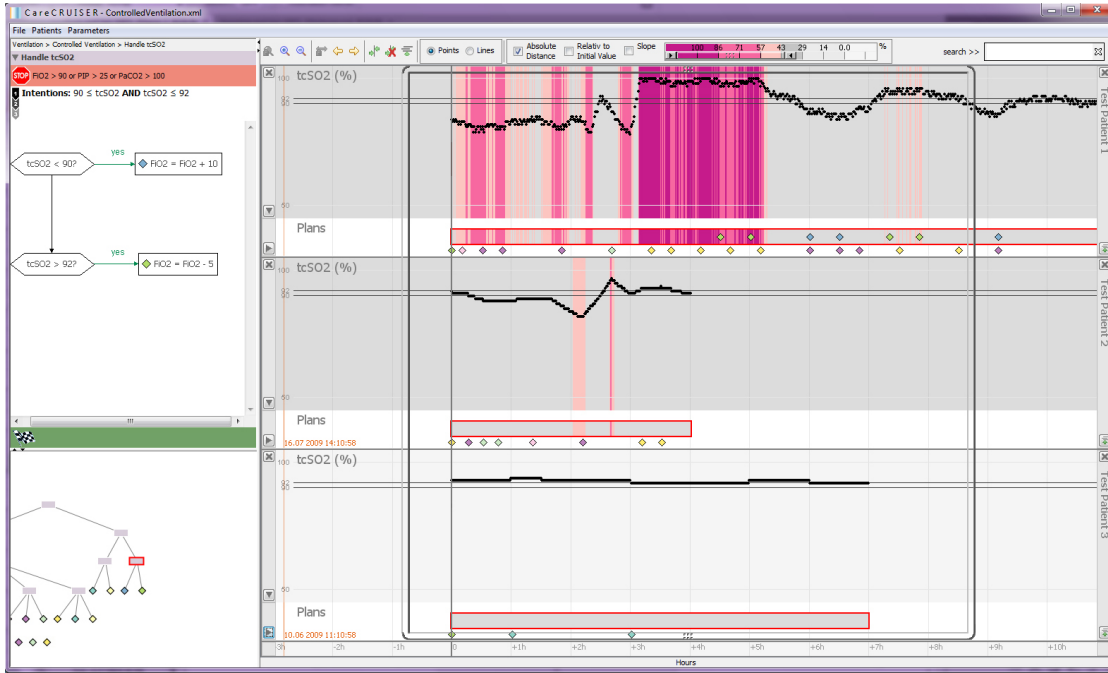


Figure 6.2: The same treatment plan has been applied to three different patients (aligned vertically). The charts and treatment plans are colored according to the color scheme of the parameter values' *distance to the intended value*. Selecting ranges of big *distances to the intended value* with the range slider draws attention to critical cases and brings out the differences between the conditions of the three patients.

plan overall succeeded in reaching the intended effects since start of application. Whereas the example above (see Section 6.1) highlights critical values, and thus provides a general overview over the conditions of a number of patients, the second example has its focus on the success or failure of the applied treatment plan in reaching the intended result. In doing so, it brings out the improvement or the deterioration relative to the parameter values at treatment plan start.

So far, this is only an enhancement of information that is also given by the line chart. However, in clinical practice it occurs frequently that more than one condition has to be fulfilled. 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 are to affect two or more parameters (e.g., 'stabilize the patient's $tcSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 mmHg'), things become more difficult: the 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 [33]:

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

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

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

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

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

According to these rules, the overall success of a plan that is 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)).

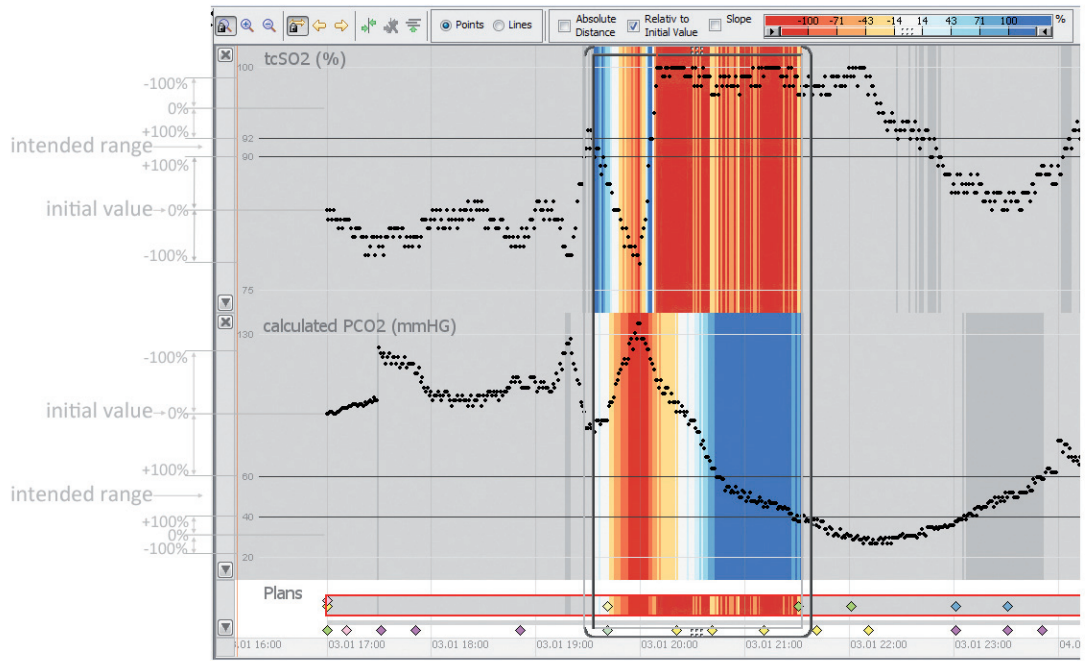


Figure 6.3: Highlighting the *progress* of parameter values *from the initial* intended value results in getting an immediate idea of success or failure of the applied treatments. In our case, the overall success of executing the selected treatment plan at a given point in time is determined by the minimum of the success values for stabilizing the patient's $tcSO_2$ and for stabilizing the patient's PCO_2 . The low success values for stabilizing the patient's $tcSO_2$ lead to a low overall success for the selected treatment plan.

Figure 6.3 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 $tcSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 mmHg. The overall success of the treatment plan execution at a given point in time is thus determined by the minimum of the success value for stabilizing the patient's $tcSO_2$ and the success value for stabilizing the patient's PCO_2 . 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, toward the end of the focus region the execution of the treatment plan has a positive effect on the patient's PCO_2 , while the patient's $tcSO_2$ values are still too high. This results in a low overall success figure of the treatment plan.

6.3 Detecting Temporal Patterns in Reactions to Clinical Actions

Reactions to clinical actions may occur right away or after some time; in the latter case, they are easily misinterpreted or overlooked altogether. These delayed reactions form typical patterns within a given temporal frame, namely patterns that reveal themselves in vertically aligned charts. They yield new insights and possibilities to optimize treatment plans; on the other hand they help to avoid erroneous procedures in medical efforts. To corroborate this statement, let's consider the following situation that we have found in the course of our testing and evaluating procedures:

A newborn has to be artificially ventilated. Accordingly, the physicians in charge lower the oxygen supplied to the patient, as the oxygen values of the patient have exceeded the acceptable optimum. As this oxygen value remains too high, another reduction in the oxygen supply is administered – which seems a reasonable thing to do in such a case. Subsequently, the oxygen value of the patient drops as expected, but it drops way below the intended value range, so that the oxygen supply has to be increased again. The visualization gives an explanation for this unexpected reaction behavior and allows the physician both to analyze and to better take into account the temporal factor in the given therapeutic process; otherwise, the treatment would lead to an oscillating reaction which can be quite harmful for the patient. This detection of the temporal pattern is shown in detail in the following five figures (Figure 6.4 to Figure 6.8) to demonstrate the versatile potential of CareCruiser's options for a stepwise interactive investigation of such patterns.

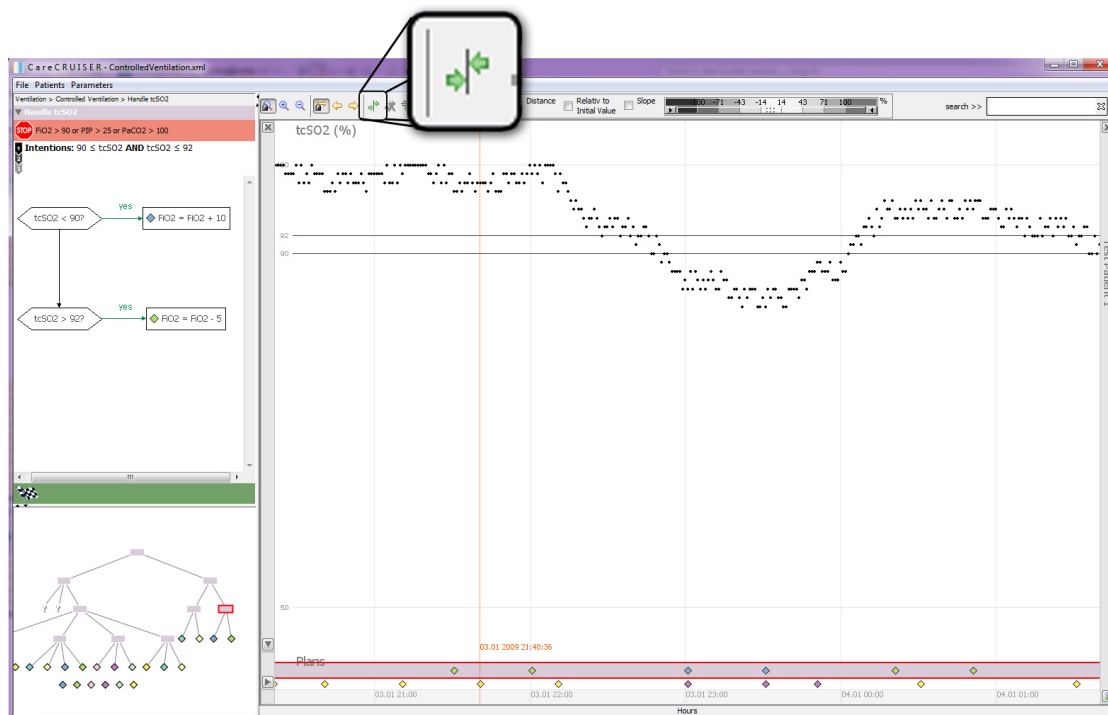


Figure 6.4: This is a standard representation of a patient's parameter values. The *Temporal View* also displays the applied treatment plan and clinical actions; logical and hierarchical representations of this treatment plan are displayed on the left (provided by CareVis [5]). Clicking the button indicated by the blow-up window and subsequently selecting a specific clinical action (diamonds at the bottom of the *Temporal View* or within the *Hierarchical View*) results in the following screen (see Figure 6.5).

In Figure 6.4 the parameter values of this artificially ventilated newborn are plotted along a horizontal timeline.

Looking at such a chart it is not trivial to grasp the context between a clinical action and its resulting parameter value. This is due to a two-fold complication in the comprehension process. Firstly, the brain attempts to get the vertical alignment right: a specific clinical action symbol in the lower part of the chart has to be brought into relation with its corresponding values in the upper part. Secondly, there is a horizontal comparison: the eye has to try to detect any possible time dependent pattern in the minute changes in the curve (because of the lack of proximity and an appropriate alignment of chart events) that should be compared [27].

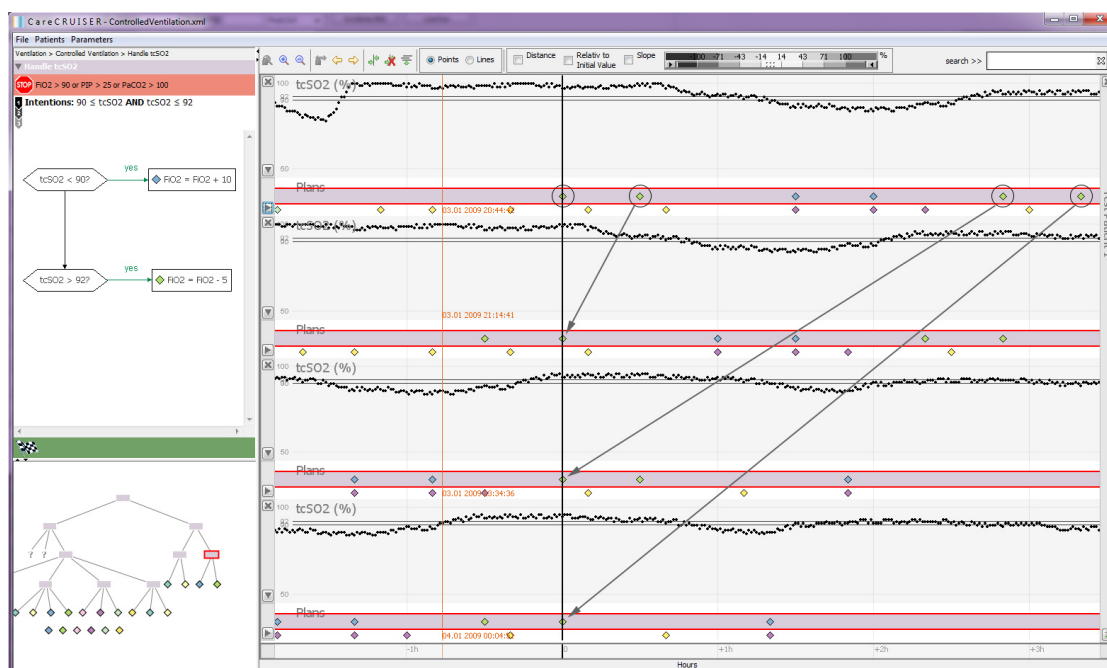


Figure 6.5: A number of identical clinical actions that have been applied to a patient at different times and the corresponding patient parameter values are aligned below each other.

To avoid this difficulty, CareCruiser, in a first step, lets you visually align all instances of the specific clinical action (i.e., the reduction of the oxygen supply) in combination with the patient parameter in question, below each other (see Figure 6.4 and Figure 6.5). The intended effect of this arrangement is, of course, to be able to more easily compare the precise development of the parameter curve after each application.

Aligning identical clinical actions one below the other makes the laborious task of comparing the resulting changes in the parameter curve much more convenient. CareCruiser allows you to do so with a simple click. The thus generated chart is shown in Figure 6.5. Either, identical actions taken with different patients can be aligned in such a way, or, identical actions applied to one patient at different times. In every such case, the eye is no longer forced to wander horizontally in an attempt to find similarities in the development of the parameter curve; instead, it can very conveniently judge this development by merely moving up and down along the same vertical line. Moving an imaginary vertical line from the time of a clinical action towards the right will always result in an alignment of corresponding parameter values (i.e., the reaction of the patient) that have the same temporal distance from this specific clinical action, which is impossible to achieve if you do not arrange the chart sections below each other.

The arrangement shown in Figure 6.5 serves as the basis for a number of additional refining procedures, among them the exploitation of color. As minor ups and downs in the parameter curve are sometimes difficult to grasp from simply looking at such a chart, CareCruiser facilitates matters by providing the *slope* coloring feature which is applied to every vertical chart instance

simultaneously (see Figure 6.6). With its help, the ups and downs – and thus, the similarities and dissimilarities – stand out clearly in every such instance.

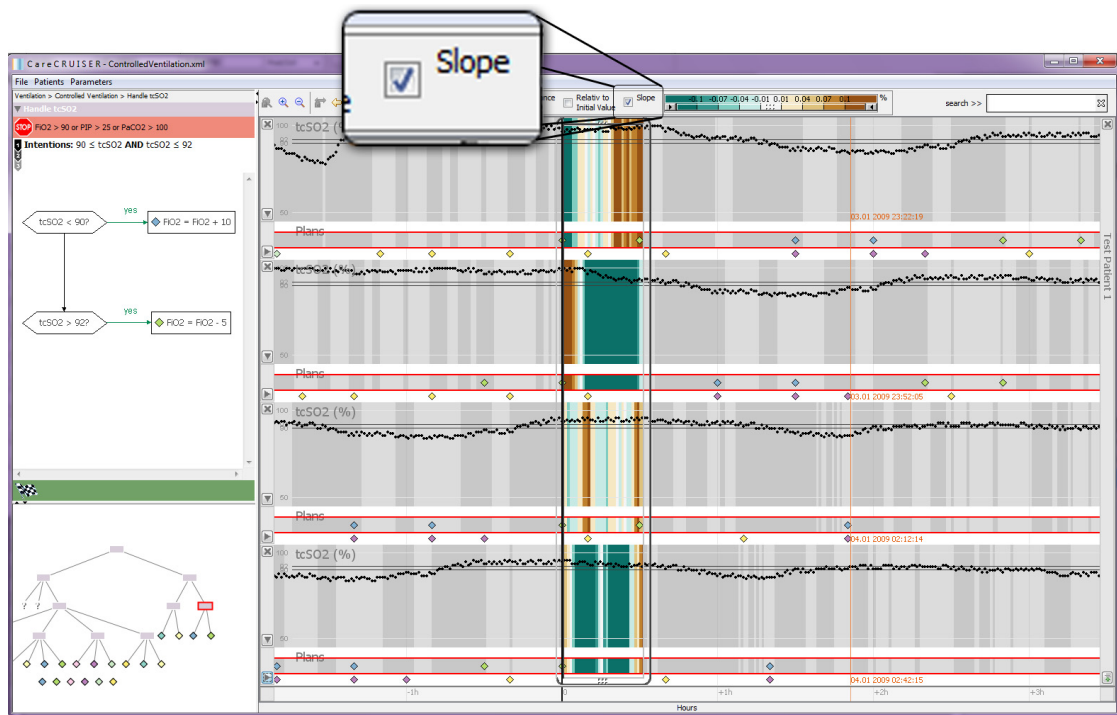


Figure 6.6: Checking the *slope* option in the toolbar on top results in the color-highlighting of positive and negative slopes (brown: positive slopes, turquoise: negative slopes) within the *Focus & Context Window*.

Since the initial intention of the clinical action was the lowering of the parameter value, we are obviously interested in the falling *slope* instances in our chart and pick them out by manipulating the range slider (see Figure 6.7). It would have been possible to color the *distance to the intended value* or the *progress from the initial value* as well, since CareCruiser provides the color highlighting of different types of information.

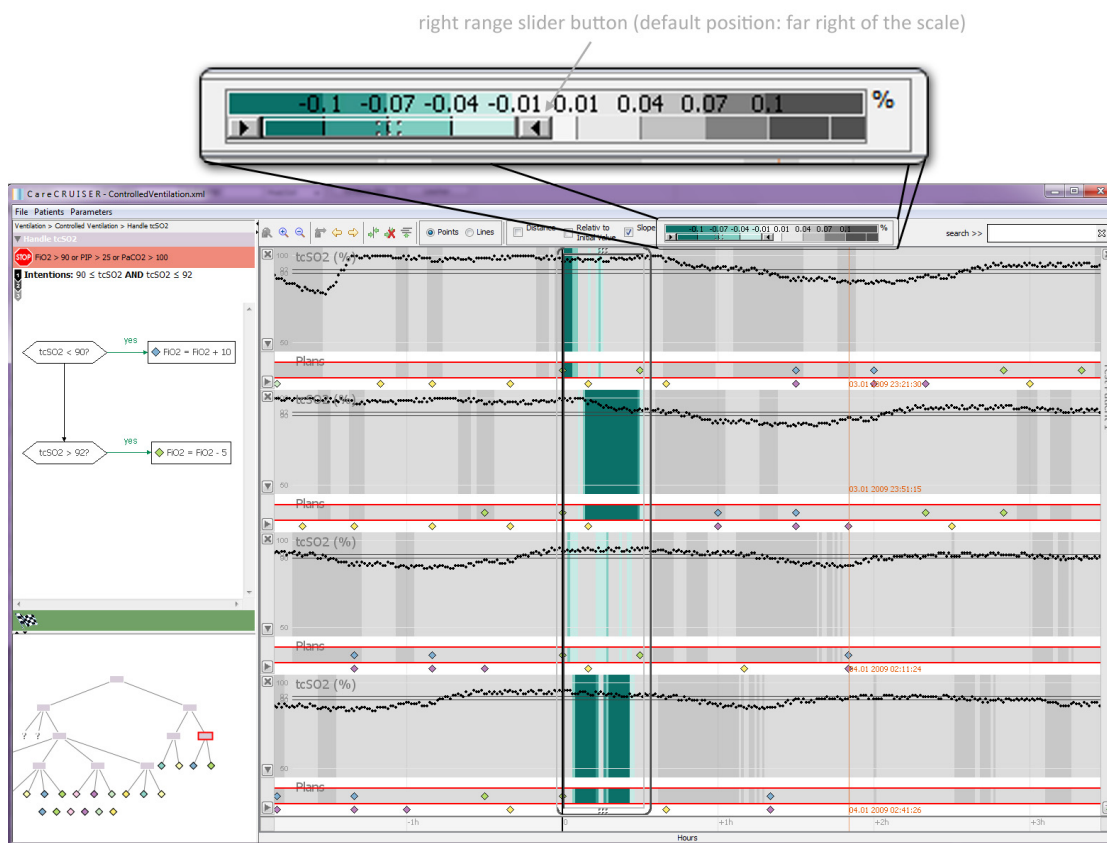


Figure 6.7: With the help of the range slider, the coloring in the *Focus & Context Window* can be reduced to events of special interest (in this screenshot the coloring has been reduced to instances of a falling *slope* by moving the right range slider button to the left).

In the given chart excerpt, it seems that the application of decreased oxygen supply sometimes has the desired effect and sometimes it does not (see Figure 6.7). However, as CareCruiser allows you to move the *Focus & Context Window* to any part of the chart, you can probe into the temporal dimension of patient reactions by tentatively sliding the *Focus & Context Window* towards the right of the applied action until – possibly – sloping effects show up that are very similar one to the other. Having found such a complete correspondence between instances of reactions to an applied action means that the action, in all likelihood, has had the desired effect – after a longer time, however, than expected (see Figure 6.8). CareCruiser helps to make such time dependent effects stand out, and thus promises to be a valuable tool in adjusting and fine-tuning clinical actions and medical treatment plans.

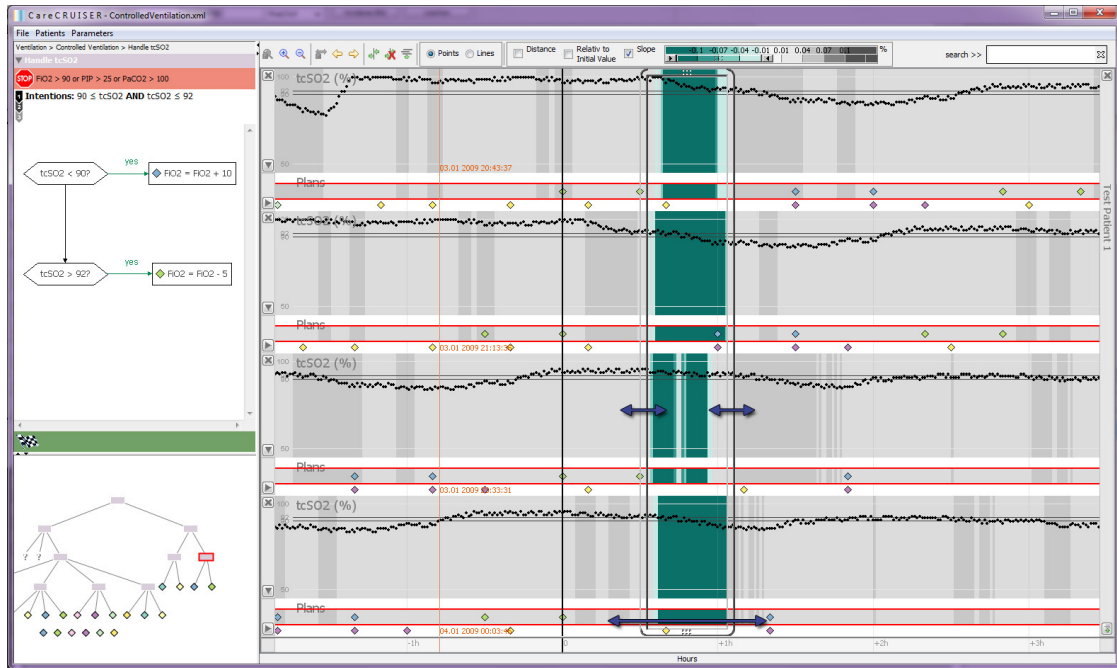


Figure 6.8: Moving the *Focus & Context Window* to different locations of the chart makes it easy to identify time and instances of identical reactions to a clinical action. The width can be adjusted as needed. The entire window can be moved horizontally in both directions.

In this section we illustrated CareCruiser’s visualization features and its potential to help physicians to investigate clinical situations in detail and thus, to improve treatment. In the next sections (Section 7.2 and Section 7.3) we describe how we evaluated our prototype to find out about strengths and weaknesses of this approach.

Part III

**Evaluation, Results, and
Transferability**

Multi-Step Evaluation Process

We took several steps to ensure both the quality and the usefulness of CareCruiser: (1) We gathered feedback from a medical expert in an early design phase to make sure that the design of the InfoVis prototype meets the actual needs of the target users. In a later stage, (2) we conducted a heuristic usability evaluation to enhance the handling of the visualization and the interaction features. Finally, (3) we tested the system with two target users (i.e., medical experts) to find out how well CareCruiser is suited to answer given research questions. For the heuristic usability evaluation as well as for the user testing, we used two different CPGs about artificial ventilation of infants together with five data sets of real patient data (four parameter sets corresponding to the first CPG and one parameter set corresponding to the second CPG).

7.1 Collaboration with a Medical Expert

In an early design phase we collaborated with a medical expert to make sure that we make the right design choices to meet his actual needs. The medical expert had formerly been working in the field of artificial ventilation for many years, and thus, he was the perfect person with whom to discuss our prototype. We presented an early version of the prototype to him and introduced all available features. While doing so, we had vivid discussions about the design choices, identified some drawbacks of the early prototype version, and also generated ideas about useful features we had not thought of originally. One observer took notes of the findings of our fruitful discussion which lasted over an hour; subsequently, we adjusted our prototypical design accordingly.

One of the most important drawbacks we identified within this session was that we used the colors blue and orange for all three coloring schemes. This had the effect that the user had to remember the active scheme very carefully; when switching between the schemes, similar colors had very different meanings which led to confusion. So, we developed the three very distinct color schemes that we are using now. Furthermore, we learned about the importance of being able to align specific events of interest, such as the beginning of a treatment or other clinical actions.

7.2 Usability Evaluation

In order to improve the usability of CareCruiser, we conducted a heuristic usability evaluation according to [30]. The reason why we decided to use this set of heuristics instead of the classical heuristics according to [62] is that this new set by Forsell and Johansson is especially tailored to the evaluation of important usability problems in InfoVis techniques. Forsell and Johansson present a best practice set of 10 heuristics out of 63 heuristics (the 10 heuristics derive from 6 earlier published heuristic sets and the numbering consists of a letter referring to the original heuristic combined with their original number):

- ‘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.
- E7. Minimal actions. Concerns workload with respect to the number of actions necessary to accomplish a goal or a task.
- 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.
- B7. Orientation and help. Functions like support to control levels of details, redo/undo of actions and representing additional information.
- 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.
- E16. Consistency. Refers to the way design choices are maintained in similar contexts, and are different when applied to different contexts.
- C6. Recognition rather than recall. The user should not have to memorize a lot of information to carry out tasks.
- E1. Prompting. Refers to all means that help to know all alternatives when several actions are possible depending on the contexts.
- 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.
- B9. Data set reduction. Concerns provided features for reducing a data set, their efficiency and ease of use.’ [30, p.203]

We conducted the study with four evaluators according to the common assumption that three to five evaluators are sufficient for a heuristic evaluation (see [45]). The two female and the two male evaluators have a degree in computer science and have gained practical and theoretical experience in the field of usability engineering.

Separate testing sessions with each evaluator are important to ensure an unbiased outcome. The evaluators were allowed to ask questions about the medical care field. In case an evaluator got stuck anywhere in the course of using the program there was always someone present to help him/her out, although this was not necessary at any time during the four testing sessions.

Each evaluator went through the program two times. The first time he/she got acquainted with the flow and the general scope of the program; the second round focused on visual and interactive interface elements with respect to the given list of heuristics (see [30]). The following tasks had to be performed:

1. Finding out how the execution of a treatment plan corresponded to a patient's condition.
2. Finding out which clinical actions were applied to the patient in the context of treatment plan execution, and when these actions were applied.
3. Finding out which effects the clinical actions had on the patient's condition.
4. Identifying critical parameter values in the course of treatment plan execution.

Any problem found was noted down with reference to the violated usability principle; the severity of the problem encountered was rated on a scale from 1 (low) to 5 (high).

7.3 User Testing

*'The purpose of computing
is insight, not numbers.'*
R. W. Hamming, 1961.

We presented the CareCruiser prototype to two physicians, physicians being the main target user group of the visualization. Both physicians work at the Vienna General Hospital – one at the Dept. of Obstetrics and Fetomaternal Medicine, and the other one at the Dept. of Child and Adolescent Psychiatry, but had formerly been dealing with the artificial ventilation of infants for many years. In this context we used a treatment plan for artificial ventilation of newborn children together with patient data of four different patients who were treated in accordance with the same treatment plan. The two physicians used CareCruiser to explore the patient data in combination with the applied treatment plans and interactively manipulated the presentation to analyze the effects of the applied clinical actions. The studies were carried out separately with each physician; in both cases the interaction with CareCruiser took about two hours. After a short introduction of the prototype and its features, the physicians used the prototype autonomously to investigate the given clinical situation. Two observers were present to document the findings.

In order to ensure both the reliability and the validity of the findings of the user testing, we collected data from different sources. This technique is known as triangulation, which is highly recommended by many researchers (see, for example, [59, 61, 99]). Two observers took notes while the medical expert was interacting with the prototype. One observer took down general observations (e.g., how the physician proceeded, which features he investigated more carefully than others, etc.), the other one took down the findings produced by the *think aloud*

method. After the interaction phase, one observer conducted a semi-structured interview with the physician (a mix of open and closed questions). To make sure all the details were captured, the interview was audio-recorded. Subsequently, the interview was transcribed into text and the transcripts were reviewed by the interviewees to make sure they agreed with the interpretation of what was said before it was analyzed.

In the first session, the test person (i.e., one of the physicians) started to make himself familiar with the logic and medical recommendations of the CPG by using the *Hierarchical View*, the *Logical View*, and the Linking & Brushing coordination between these views. Then he continued to investigate the patient data and the actually applied clinical actions in the *Temporal View*. The next step was to modify the set of parameters displayed in the *Temporal View*. He reconstructed why the clinical actions had been applied in the given situations, and identified some actions which had not been applied according to the CPG recommendations; this led to the identification of situations in which the patient was not optimally treated. He used the highlighting of the parameter's *distance to the intended value* to get an impression of the overall clinical situation during treatment plan application (i.e., the identification of very critical periods and the overall improvement of the patient's condition during treatment). Moreover, he used the coloring of the parameter's *progress from the initial value* to judge to what extent the treatment plan succeeded in reaching the specified goals (i.e., in affecting the patient's parameters according to the specified intentions of the treatment plan). When examining the effects of single clinical actions in more detail, he used the feature of *aligning actions vertically* in combination with the coloring of the parameter curve *slope*. In doing so, he verified if the clinical actions resulted in raising or lowering the values of the patient parameters in question.

Quotations in this context:

'The coloring [of the parameter *distance to the intended value*] seems to improve the visibility of the patient's condition.'¹

'The visualization really helps to reconstruct which clinical actions have been applied as well as to find out why they have been applied in given clinical situations.'²

In the second session, the other physician immediately started reconstructing the reasons for applying specific clinical actions in given situations. While doing so, he gradually adapted the set of parameters displayed in the *Temporal View* according to his needs. He highlighted the *distance* of the parameter values *to the intended value* range to get an overall impression of the situation. For each clinical action that had been applied he decided – with respect to the given clinical situation (i.e., the patient's condition) and with respect to the recommendations of the CPG – if it corresponded to a best possible treatment of the patient. In a second step he aligned selected clinical actions vertically and used the *color-highlighting* of the parameter curve *slope* to check the effects of clinical actions. For a more detailed examination of the actual effects of a specific clinical action that was supposed to lower specific parameter values, he filtered for negative slopes and used the *Focus & Context Window* to inspect a short span of the time that followed the application of the action. The physician reasoned that the patient did not receive the

¹translated from German by the authors: 'Die Farben scheinen die Übersicht über den Patientenzustand zu verbessern.'

²translated from German by the authors: 'Bei dieser Visualisierung ist sehr schön nachvollziehbar, welche Aktionen warum gesetzt wurden.'

best possible treatment. Interestingly enough, he agreed that the findings made with the help of CareCruiser illustrated that the effect of a specific clinical action always happened with a certain delay, which seems to have been the reason why, in the case of the given patient data, the action was applied two times in a row. This consequently led to an excessive drop of the parameter values below the intended value range, indicating a less than ideal patient treatment (compare the example we describe in detail in Section 6.3). The physician also agreed that the interactive InfoVis method as offered by CareCruiser helps to easily pinpoint these treatment deficiencies and thus to adapt the treatment plans accordingly, in order to avoid these deficiencies.

Quotations in this context:

‘This is a very feasible solution to analyze this kinds of data, with interesting features.’³

‘The strengths of the visualization are that it is clear, easy to work with [...] - intuitive, very nice.’⁴

‘I believe [with the help of this visualization method] one can well demonstrate a clinical situation; e.g., [with its help] one can explain the situation to one’s students.’⁵

After the actual evaluation sessions we carefully sorted and analyzed the derived data in order to generate meaning and to draw conclusions. An outline of the results is given in the following section.

³translated from German by the authors: ‘Das ist eine sehr praktikable Lösung, diese Daten zu analysieren, mit interessanten Features.’

⁴translated from German by the authors: ‘Die Stärken sehe ich darin, dass die Visualisierung übersichtlich ist, dass sie handlich ist [...] - intuitiv, wunderschön.’

⁵translated from German by the authors: ‘Ich glaube, dass man die klinische Situation gut demonstrieren kann. Dass man z.B. seinen Schülern die Situation erklären kann [...]’

Outcome

A detailed description of the outcome of the usability evaluation and the user testing.

8.1 Usability Evaluation

32 usability problems were encountered in the course of the usability evaluation. Most frequent among them were the two usability principles of ‘orientation and help’ and ‘consistency’. ‘Orientation and help’ stands for controlling levels of details, redo and undo of actions, and giving additional information; ‘consistency’ refers to how well a design choice was kept up in similar contexts (and its opposite: how well a design choice was contrasted with another in differing contexts) (see [30]). Table 8.1 and Table 8.2 show the number of problems concerning the violated usability principle and the number of problems with respect to the severity assigned by the evaluators.

Principle	B7	E16	B5	E11	D10	E7	B9
Problems	12	7	3	3	3	3	1

Table 8.1: Usability problems. The Distribution of usability problems with respect to the violated principle.

Average Severity	1	2	3	4	5
Problems	6	7	10	6	3

Table 8.2: Usability problems. The Distribution of usability problems according to their average severity (5 representing the highest severity).

Most of the problems found are not severe and could be fixed immediately. There have been found, however, a couple of problems that will require some additional effort in order to be repaired:

- One of them is the lack of a possibility to align the grid with the clinical actions.
- Another one arises in the following situation: In case the user wants to select a single instance of a clinical action that was applied multiple times, all instances of this action are highlighted, so that there is no difference to be seen between the one instance the user has in mind and all the other instances.
- As far as the undo/redo capabilities of the program are concerned, the evaluators would have liked to have a history of view modifications and to be able to jump back and forth between views, including the optional skipping of any number of modification steps. This seems to be a very handy feature that warrants its implementation.

We take all of the found problems seriously and – if possible – we shall correct them in the near future.

In the following section we take a closer look at the findings of the user testing.

8.2 User Testing

Both physicians reacted in a very positive way to using CareCruiser, as is evident from the quotations above which were picked to corroborate our findings (the complete interviews can be found in Appendix B). It took both of them very little time to learn how to use the system. After approximately 15 minutes they were able to navigate between views, add and remove patient parameters, apply different kinds of color-highlighting, and align clinical actions. While interacting with CareCruiser, they came across unexpected treatment choices, modified the view to explore the situation in more detail, and appraised the quality of the treatment. In addition, they found unexpected patient reactions to clinical actions; one physician formulated hypotheses on how to optimize the treatment plan in accordance with these findings. Both physicians eagerly used CareCruiser's interactive features to explore the given clinical situation. When being asked about shortcomings of the interactive InfoVis prototype, they suggested the following additions and changes:

- adding labels to the tree in the *Hierarchical View* and the clinical actions in the *Temporal View*;
- indicating if a ventilation adjustment should be increased or diminished (clinical action) in addition to giving a formula for the calculation of the actual value;
- enlarging the labels of the parameter scales in the *Temporal View*;
- using thicker lines to indicate the intended value ranges in the *Temporal View*; swapping the position of the *Logical View* and the *Hierarchical View*;

- adding the possibility to flexibly define a reference point on the timescale and to highlight the bettering or worsening of the patient's condition after that point (similar to highlighting the *progress from the initial value*).

After the interaction phase, we conducted semi-structured interviews with 13 questions (the complete list of these questions and transcribed answers can be found in Appendix B), which can be summed up in the following three main questions:

1. Does CareCruiser provide appropriate information in an intuitive way?
2. Does CareCruiser help to judge if a treatment has the intended effect on a patient's condition?
3. Does CareCruiser help to detect unexpected effects of clinical actions and thus help to optimize treatment choices?

When being asked if CareCruiser helps to faster form an opinion about a clinical situation, one physician said that in case of an emergency the standard kind of acoustic alarm would do a better job. (Replacing acoustic alarms in critical situations has never been the intention of our prototype, though.) Both physicians gave positive answers to all the questions relating to the three main categories mentioned above:

1. They said that CareCruiser facilitates capturing the course (i.e., positive and negative effects) of an applied treatment. Quotation in this context: 'What is very nice is that here you can see on a temporal axis – beside changes in the patient parameters – all the actions that have been applied. Thinking about it, one can, of course, see in particular the effects of actions taken before, which means that you can scrutinize anything that one has tried out before and realize what has already been successful.' ¹
2. They affirmed that CareCruiser provides all information (about the patient and the treatment plan) necessary to assess the clinical situation. Quotation in this context: Interview question: 'Is CareCruiser a suitable instrument to represent all the necessary data to be able to assess the health state of the patient in an adequate way?' Answer: 'Yes, the patient data suffice absolutely.'
Interview question: 'Is CareCruiser a suitable instrument to represent all the necessary information about the applied therapy in an adequate way?' Answer: 'The information [note: about the applied therapy] suffices as well.' ²

¹translated from German by the authors: 'Was hier auch sehr schön ist, ist, dass man neben den Veränderungen der Patientenparameter auch alle Maßnahmen, die gesetzt wurden, in der Zeitachse sehen kann. Insbesondere – wenn ich jetzt näher darüber nachdenke – sieht man natürlich auch die Effekte von Maßnahmen, die man vorher schon gesetzt hat, d. h. man kann sich auch anschauen, was man vorher schon ausprobiert hat und was schon erfolgreich war.'

²translated from German by the authors: interview question: 'Ist CareCruiser geeignet alle nötigen Informationen um den Patientenzustand einschätzen zu können in adäquater Form darzustellen?' Answer: 'Ja, die Patientendaten reichen vollkommen aus.'

Interview question: 'Ist CareCruiser geeignet alle nötigen Informationen der angewandten Therapie in adäquater Form darzustellen?' Answer: 'Die Informationen [Anmerkung: über die angewandte Therapie] reichen auch aus.'

3. They stated that CareCruiser helps to detect unexpected patient reactions to applied clinical actions. Quotation in this context: Interview question: ‘Do you think CareCruiser can help to detect unknown/hidden reactions to clinical actions?’ Answer: ‘Yes, definitely’³. Additional quotation in this context: ‘In particular, when investigating a longer period of time, the various actions and their effects are clearly discernible. I really cannot see how to derive this information in any other way. In my opinion, this kind of visualization is the only feasible way of bringing to light the effects of the applied clinical actions.’⁴

Generally speaking, they were sure that the insights gained with CareCruiser help to optimize treatment. One physician stated that the interactive InfoVis prototype might also be of great value when a patient is handed over to another physician (e.g., a shift handover), since it provides useful information about which clinical actions had positive effects on the patient’s condition and which had not. General findings of the study infer that the interaction techniques, in spite of their complex nature, are well understood and accepted. New possibilities to explore the data at hand are welcome if they lead to a better understanding of the situation, and the color-coding and highlighting of derived values in combination with raw data values helps to bring into focus specific events that are easily overlooked otherwise.

³translated from German by the authors: interview question: ‘Glauben Sie, dass CareCruiser helfen kann, nicht bekannte / nicht offensichtliche Reaktionen auf klinische Aktionen zu erkennen?’ Answer: ‘Ja, absolut.’

⁴translated from German by the authors: ‘Gerade wenn man einen längeren Zeitraum betrachtet, sieht man die einzelnen Aktionen und was sie bewirken. Ich wüsste nicht, wie man das sonst in irgendeiner Art und Weise ablesen könnte. [...] Ich glaube, diese Visualisierung ist überhaupt die einzige Methode, diese Effekte von den gesetzten klinischen Aktionen überhaupt sichtbar zu machen.’

Transferring to Other Domains

It is our intention to extend our research to other areas as well, where our visual encodings might be of advantage. Our InfoVis prototype can be put to use in any domain in which it is desirable to investigate the effects of actions, and, in particular, to trace back middle- and long-term effects to their causes. This may hold true for the field of economics (e.g., stock market proceedings and their effects on currency rates), for the field of public relations (such as the effects of press reports and publicity measures on the approval rates for politicians), the academic field (such as the effects of introducing new structural or curricular elements in education on the quantity and quality of graduates and/or scientific output in different countries), a.s.o.

9.0.1 Approval Ratings

The following simulation attempts to illustrate the potential of our interactive InfoVis technique in the field of politics and in the field of market research. We have assumed a specific situation that is being influenced by certain actions, therefore, we have created artificial data of a sort that could conceivably occur in this context. A politician depends very much on the approval of those whom he/she represents. This degree of acceptance by the public is a volatile phenomenon, steered both by facts such as unemployment, living costs, etc., and by public relations, such as television appearances and the print media reports. Of course, a politician can always ask a market institute to find out how much approval or disapproval concerning himself/herself there is. However, this is a very rudimentary tool that informs him/her about the momentary state of things only; the visualization method developed for CareCruiser can help him/her to find out about the dynamics that have led to the static figures that he/she has to rely on otherwise.

Suppose a politician leads a private life of a rather scandalous kind and papers would eagerly produce reports on it. In Figure 9.1 we have color-highlighted the slope of the approval rating curve and manipulated the range slider to highlight only the negative falling of the curve (i.e., the approval ratings), while the negative press reports are represented by the diamonds at the bottom. Two effects can be imagined: either each of these reports results in a temporary disapproval of the public followed after some time by a period of forgiveness ('we are all human beings') and

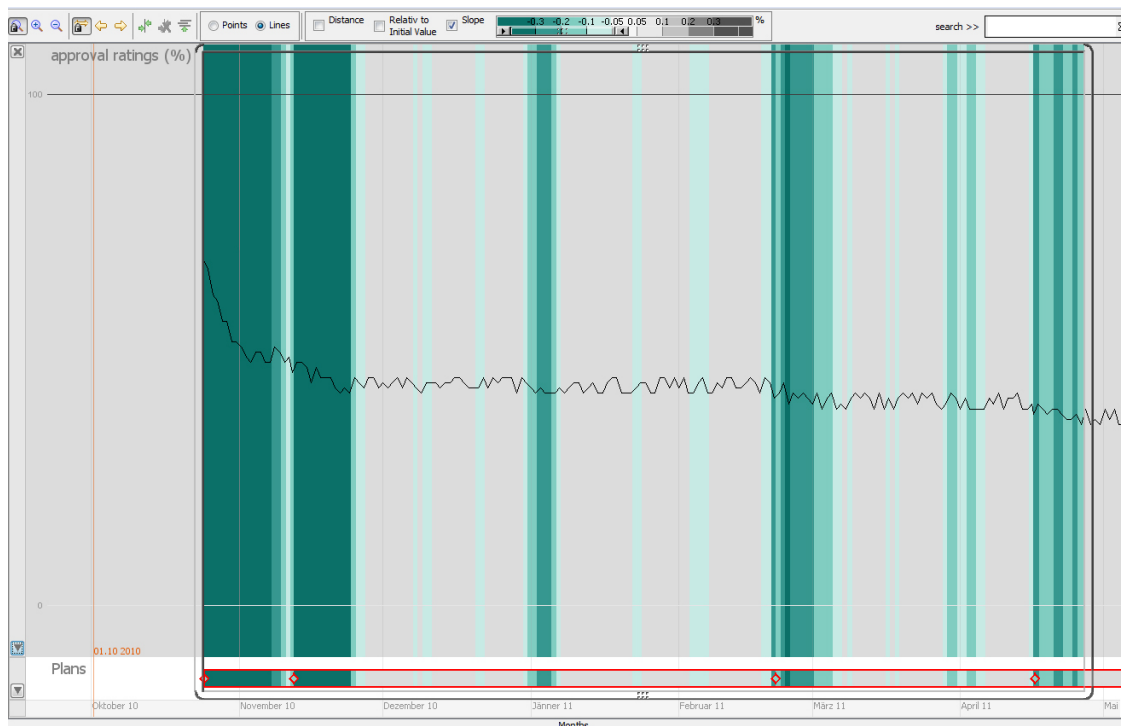


Figure 9.1: Approval ratings. In this case the ‘events’ (the diamonds at the bottom) are negative press reports and the curve shows the approval ratings. We can see the curve has a downwards trend but the actual effects of these reports are not clear by now.

even silent approval or the curve does not recover. The first case is simulated in Figure 9.2, where the use of color helps to immediately assess the situation. Similar to the example given in Section 6.3, we have aligned all negative press reports at a vertical axis and moved the *Focus & Context Window* to the time period immediately after the negative press reports came out. Here, the color allows to identify subtle patterns in a way that could be achieved only with difficulty otherwise. In the case assumed in Figure 9.2 the cone shape of the color pattern and its fading-out towards its lower portions indicate that these negative newspaper reports do not really harm the politician; this can be deduced at a very early stage by precisely comparing the effects of a couple of these reports. The politician would then not have to concentrate on any sort of cover up operation or beginning a new kind of life which would certainly be advisable for him to do if the fade-out and tapering of the color pattern of the curve did not occur.

In case the curve starts falling and does not recover – instead, every newspaper report results in a steeper inclination of the curve – a number of *slope* factors could be deduced relating to the urgency of actions the politician has to undertake in order to stop this trend. It can even be assumed that there exists a *slope* factor that indicates to the politician with absolute certainty that his/her case is lost. We may assume that finding out as early as possible about such a lost case would be of decisive importance, as it would enable the politician to leave his/her post in

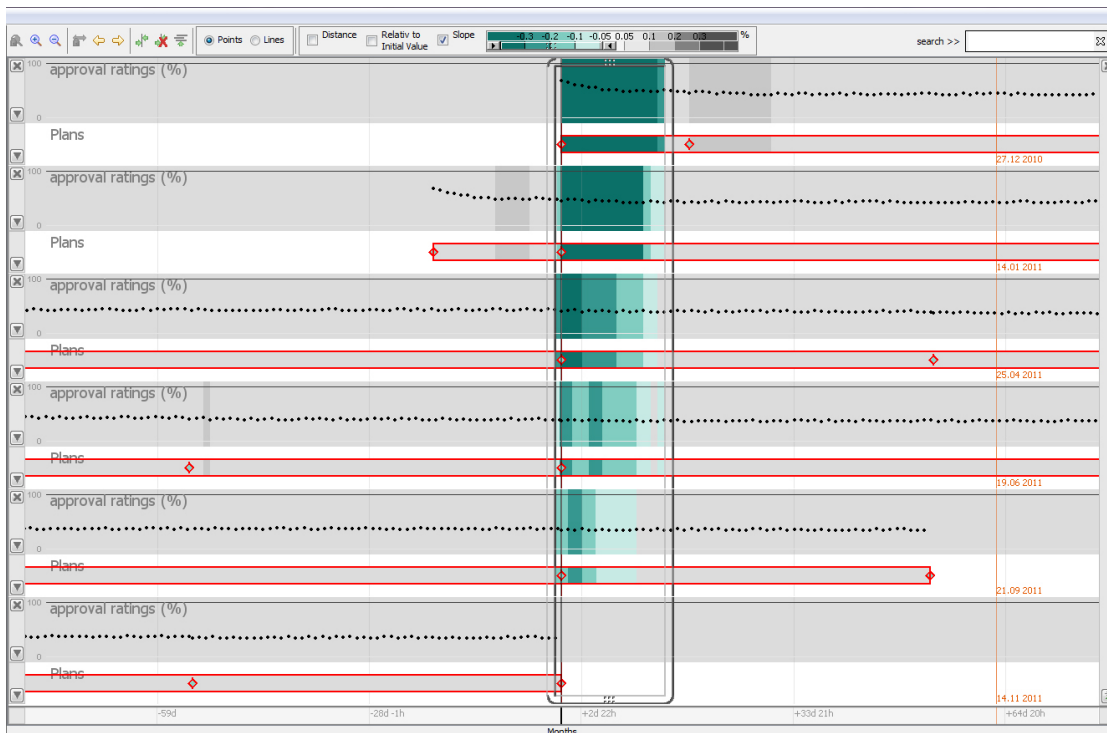


Figure 9.2: Approval ratings. In order to investigate the effects of the negative press reports on the approval ratings, we have aligned these negative reports and color-highlighted the *slope* of the curve. We can see that these reports result in a temporary disapproval but their negative effects diminish over time and it can be assumed that they don't really harm the politician in the long run.

an orderly manner (installing a successor, finding a new position, etc.) instead of simply being eliminated from his position.

9.0.2 Sales Figures

The second example deals with market products. It goes without saying that here, too, the time factor is of essential importance. Basically, we are confronted with a scenario similar to the first one, the only difference being that in the second scenario it is the sales figures that count (and not the approval rates). However, we use a different aspect of our interactive visualization method, as we are interested in slightly different aspects of the development of these marketing figures.

Here we assume that a new product appears on the market, is well accepted and consequently is produced in ever growing numbers. To keep the acceptance level high, adequate PR actions are taken. After a certain time, the sales figures grow more slowly because the market becomes saturated. Depending on the sort of product, two courses of action are advisable: Firstly, the producer knows that in order to keep the sales figures on this level, he/she must keep promoting it; at the same time he/she must adjust production to meet this saturation level. Secondly, if

the producer intends to expand, he/she would have to create something new: this could be a product that has new features, or it could be a new, more aggressive way of promoting the old product. Our visualization method displays the probable development of the sales figures at any given moment in time and thus would enable the producer to finetune the deployment of his/her means. In Figure 9.3 we can see that the sales figures rose rapidly when the product appeared on the market. After some initial oscillation the sales figures even out at a given saturation level while PR actions are taken in frequent intervals. When the producer starts to take these PR actions more sparingly, a falling trend of the sales figures can be identified. In Figure 9.3 we have applied the *distance to the intended value* color scheme. In this case, the intended value range indicates the desired acceptance level of the product. The color-highlighting emphasizes the negative effect of the more sparingly taken PR actions. Besides emphasizing the information that is given by the curve, this visualization could help to monitor and compare at one glance a couple of products on the market (similar to Figure 6.2).

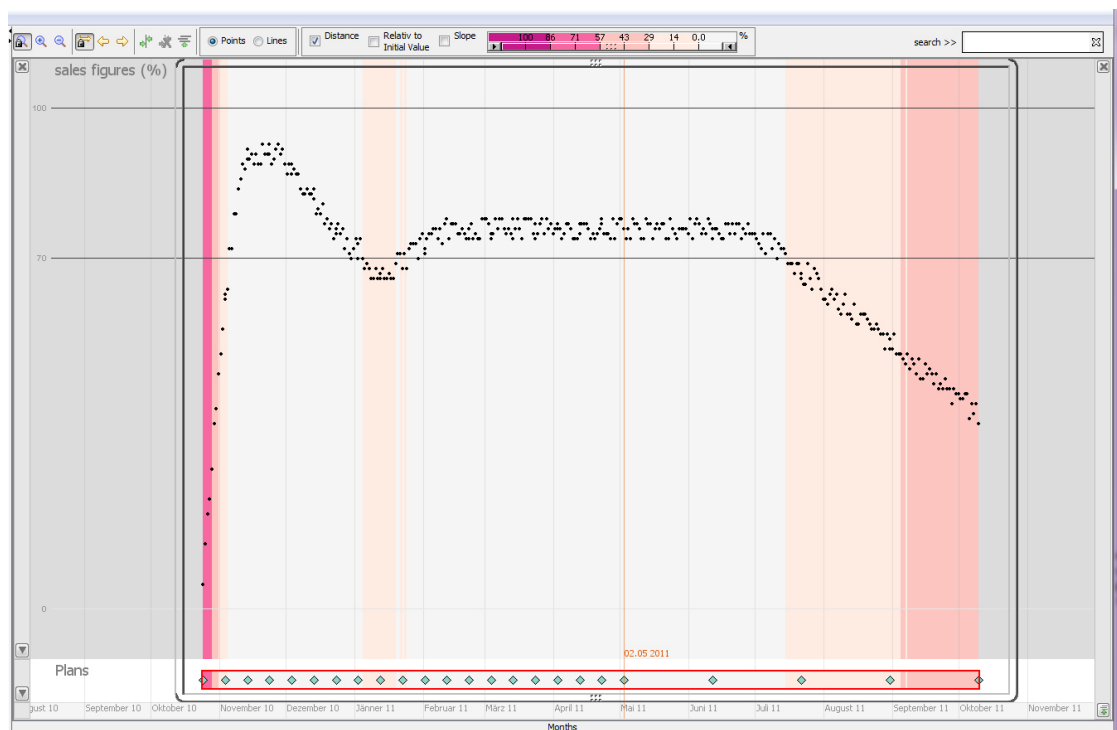


Figure 9.3: Sale figures. The ‘events’ are PR actions and the curve shows the sales figures of a new product. The new product is very well accepted and its sale figures rapidly grow. After some oscillation at the beginning the sales figures even out at a given saturation level while PR actions are taken constantly. When the PR actions are taken more sparingly we can identify a negative effect on the sales figures.

Part IV

Conclusion & Future Directions

In this final chapter we repeat the most important aspects of our work: we give a short summary of what has been done, we lay out how we have answered our research questions, we list our publications and the main contributions we have made to the scientific field of InfoVis, and we sketch our further research directions.

9.1 Summary

The steps that we have taken consist of

1. A literature research as to the information needs of patients and clinical personnel,
2. An in-depth analysis of the data characteristics,
3. Finding out in how far existing InfoVis methods can cope with the data at hand,
4. Developing a prototype in order to overcome the shortcomings of existing approaches, and which adds the requirements and design criteria resulting from the data and task analysis, and
5. Evaluating this prototype.

ad 1: There exist several papers dealing with various aspects of the things that are likely to be of interest to patients, nurses, and/or physicians. We have conducted a **literature research** pertaining to these papers, summarized the findings, and categorized the kinds of **information needs** that have become evident in the course of this research [39].

ad 2: We have looked into **what sort of data** there are when dealing with clinical treatment plans and patient data. We have found out that we have to deal with (a) hierarchically structured treatment plans that contain a strong temporal aspect; moreover, these treatment plans contain logical sequences that have to be made visible, too.

We also have to deal with (b) patient data that underlie strong temporal constraints. Patient data contain both quantitative values and qualitative values. Consideration has to be paid to the fact that the patient data may be of a high frequency nature (i.e., they contain many measurements taken within a relatively short period of time) and the opposite, namely data of a low frequency nature as well.

We realized that the combination and coordination of these distinctly different data types is of great importance in efficiently supporting the various medical tasks at hand (i.e., the in-depth investigation of effects of treatment plans and clinical actions on a patient's condition).

ad 3: Among the numerous, often very sophisticated, **existing visualization approaches** that we have looked into, we have found quite a few that were capable of coping with certain aspects of our data in a satisfactory way; however, none of them lent itself readily to fulfilling all of the requirements that have to be met.

ad 4: As a consequence of our findings, we have developed an interactive **visualization prototype** in order to make up for the deficiencies of other programs. Our interactive InfoVis prototype provides features to investigate the effects of treatment plans and clinical actions on the patient's condition in detail. In particular, our prototype provides means (1) to assess the success or failure of applied medical treatment plans as far as reaching their intended goals goes, (2) to investigate the effects of applied clinical actions (e.g., the administration of a certain drug) on selected patient parameters, and (3) to appraise the quality of the treatment and, if possible, to optimize the treatment accordingly (be it that the treatment is being tailored to the individual patient or that the treatment plan needs revision of a general sort).

ad 5: Besides **guiding our design process** by gathering feedback from a medical expert on our design choices, we took the following actions in order to **evaluate** the quality of our prototype. We have conducted a usability study using a usability heuristic that is especially tailored to the InfoVis domain [30]. We conducted our study with four evaluators with practical and theoretical experience in the field of usability engineering. This was followed by a user testing with the help of two physicians.

A number of usability problems came to light in the course of the usability evaluation and gave important hints as to what aspects of the prototype can still be improved.

The user testings have yielded satisfactory results; working with the prototype was received favorably by the test participants and it turned out that it was easy for them to work with the features provided by our interactive InfoVis prototype. They made helpful suggestions as to how this prototype can still be improved. In the course of the testing procedures, they hit upon unexpected treatment procedures in our dataset and also found a hitherto overlooked pattern in patient reactions to a certain clinical action. In particular, they noticed a less than optimum patient treatment, as the clinical action of raising the oxygen supply in the course of artificial ventilation led to an unexpected delay in the patient reaction to this oxygen supply which in turn made the physician raise the oxygen supply even further. This led to a very unstable and dangerous oscillation of the patient's pressure of carbon dioxide (PCO_2). This showed that our prototype is useful in identifying these kinds of time-dependent effects and their patterns (by means of aligning the medical data with respect to specific events in the treatment, such as changes in the oxygen supply, and, moreover, color-enhancing the effected changes in the patient's parameters), which easily escape one's attention, otherwise. Upon coming across these unexpected findings, they interactively modified the view to investigate these phenomena in more detail, which, after all, was the original intention of designing our prototype. Thus, new insights into this kinds of medical data can be gained.

9.2 Answers to Research Questions

After this summarization of our work, consisting of the problem analysis, the design and development of a prototype, and the testing procedures, we outline a brief overview over the questions that have led to this work and over the answers we provide.

Main question:

- Can we support a **detailed analysis of - and an adequate reaction to - the interdependencies** of a patient's health condition and the applied treatment (plan) with Information Visualization?
- Answer: Yes. To do so, we have developed an interactive InfoVis prototype and tested it with physicians. The physicians used the prototype to analyze a given clinical situation and formulated suggestions as to how to improve the treatment. Which proves that our interactive visualization is well suited to support these tasks.

Sub questions:

- How can the **success or failure** of an applied treatment plan in terms of a **positive or negative change** in a patient's condition as a reaction to the applied treatment plan adequately be visualized?
- Answer: To enhance the visibility of the patient reactions to an applied treatment plan, we have developed three different methods, using color to bring out any changes of interest in patient conditions:
 - (a) There is the *progress from initial value* color scheme, which highlights in how far the patient's condition has changed since the beginning of the treatment. This could be a bettering of the patient's condition compared to the patient's condition before the treatment plan was applied, or a negative change, respectively. This color scheme visually encodes the success of the treatment in a direct way using different shades of blue for positive developments and shades of red for negative developments.
 - (b) There is the *distance to the intended value* color scheme, which highlights the critical values of a patient. This fulfills a twofold purpose: firstly, it shows the situation of a patient's condition at any given moment of treatment plan application, and secondly, it also provides an integral view of the tendency of the patient's condition.
 - (c) There is the *slope* color scheme, which highlights the rise and fall of a patient's parameter. This is especially useful to check in how far a patient reacts to an applied clinical action.

Again the feedback of the physicians gave proof to us that the different color schemes were not only understood by them but also applied to investigate the details of the clinical situation for which the color schemes were developed.

- How can the visualization support the **comparison** of the effects of the **same treatment** on the conditions of **a number of patients**?
- Answer: In our prototype, a number of patient graphs can be compared with one another in the same window. This, in combination with the appropriate color scheme, leads to an immediate identification of similarities and dissimilarities in the reactions of a group of patients to the same treatment plan instances.

- Which means are suited to **detect temporal patterns** of the effects of clinical actions?
- Answer: The detection of temporal patterns is of central importance in our work. Thus, we have devised a combination of means to make them stand out: firstly, we provide the means to align actions or plans. Secondly, we use color to enhance the visibility of minute changes in the patient parameter curve. Thirdly, further enhancement is provided by our means to filter the color-highlighting of curve events which helps to concentrate on one aspect of particular interest. Last, but not least, there is the *Focus & Context Window* to show only a desired time span in full color, leaving everything else grayed out. With the help of these means it was possible to detect a special pattern of patient reaction to a given treatment (as described in Section 6.3).

9.3 Publications

In the course of this work we have published some of our findings at international conferences and in journals.

We published the findings of the preliminary literature research about the information needs of medical personnel and patients in the Journal of Evaluation in Clinical Practice [39]. We provide categorized information about information needs of the different user groups of computer executable CPGs. Not only do we cover the information needs arising in the course of CPG execution (i.e., information needs of physicians, nurses, nurse practitioners, and patients), but we also give detailed information as to the tasks of guideline modelers.

In order to ensure the quality of the research subject and to discuss the initial design ideas with InfoVis experts, we presented a paper about the PhD concept proposal at the IEEE VisWeek 2009 Doctoral Colloquium [35]. This doctoral colloquium was a great help in refining the research questions and the scope of the project. Moreover, we could thoroughly discuss our design ideas and gather feedback from experts which led to a better understanding of InfoVis methods and how they can best be applied to assist user tasks such as browsing the data and finding patterns.

At the IEEE Pacific Visualization Symposium, we presented our complete prototype together with the findings of the heuristic usability evaluation [36]. In this paper, we focused on the detailed description of the CareCruiser prototype, the visual encodings, and the prototype's interactive features. In addition, we illustrated how CareCruiser can be used to give a picture of the patient's condition, to compare different patients, and to interactively re-arrange the given graphs in order to identify any patterns of effects of applied clinical actions.

After conducting testing sessions with target end-users, we presented our findings at the BCS Conference on Human Computer Interaction [37]. In this paper, we focused on the findings of this user testing. We presented information about how long it took two physicians to learn how to use the system, in which way they actually used CareCruiser to investigate the data, what difficulties they had in handling the prototype, and how using CareCruiser led to the formulation of hypotheses about treatment choices and about effects of clinical actions.

In the first International Workshop on Theory and Applications of Process Visualization, we published a novel InfoVis method to represent CPGs with special respect to their nesting [77]. In particular, this method was designed to facilitate the modeling process of the complex hierarchies of CPGs in the Asbru representation languages [58, 78]. The visualization of CPGs is integrated into an editing environment, which has the advantage of giving an immediate visual feedback of changes.

Moreover, we will present a position paper [3] at the 2012 Workshop on Visual Analytics in Healthcare which is held in conjunction with the IEEE VisWeek 2012. In this paper we present six challenges for Visual Analytics in healthcare. We discuss CareCruiser as a possible answer to one of these challenges, in particular, to the challenge of intertwining a patient's condition with the treatment process.

Our further research led us – amongst other directions – to the problem of data quality. In particular, we probed into the quality of time-oriented data, which is a big issue in the medical domain. At the International Cross Domain Conference and Workshop on Availability, Reliability and Security, we presented a taxonomy of quality problems of time-oriented data [38]. Time has very special characteristics and adds much complexity to the data which was hardly considered in general taxonomies of data quality issues. Since the quality of the results of any data processing task highly depends on the quality of the input data, a data cleansing step would be of great importance before providing a visualization of the medical data (e.g., with CareCruiser).

9.4 Main Contributions

Originally, we were confronted with the problem of adequately visualizing the effects of medical treatment plans on patients. To this end, we conducted a **literature research** in order to find out about existing approaches; among the many existing programs, it is CareVis that stands out as particularly well suited to visualize important aspects of a treatment plan execution. The large majority of the other programs are well suited for specific purposes, but they do not sufficiently meet our requirements. In addition to conducting this research of existing InfoVis approaches, we also probed into the manifold information needs of the groups of persons that are involved in executing and undergoing medical treatments. We collected this information from various scientific publications, combined the individual chunks of information, and organized and categorized them to form a **comprehensive knowledge basis** for our work and for further research. To make up for the lack of a suitable method for the problem we wanted to tackle, we designed a **new interactive InfoVis approach**, developed a **prototype**, and **qualitatively evaluated** its usefulness.

A great boon to our work was the existence of CareVis, which provided a suitable fundament for our diverse extensions. **Integrating new methods** into this existing InfoVis prototype has the advantage of allowing us to **keep valuable and useful parts** of CareVis. We have deliberately kept its detailed visualization of treatment plans as well as the visualization of patient parameters by line charts, which is the method that physicians are used to. To this existing groundwork we have added a number of features with the aim of **making treatment aspects visible that otherwise tend to be overlooked** – which has been demonstrated by our evaluation.

Thus, it was the **new combination of suitable interactive means** that has turned out to make a difference and appeared to be a successful approach to fill the gaps of existing InfoVis approaches. This combination is made up of a careful, sophisticated mix of

- color
- focusing features
- filtering, and
- re-arranging graphs of selected time spans

and has resulted in an original InfoVis method to **analyze and optimize existing patient treatment**. We believe that these visual encodings and interactive features are of great advantage in various domains dealing with the investigation of temporal effects of events on quantitative parameter values, as we sketched in Chapter 9 how this methods could be utilized in other domains.

9.5 Further Work

Besides debugging and fine-tuning our prototype, there are a number of points in connection with our work that lend themselves readily to potential future fields of research.

Finetuning of the Prototpe

As described in Section 8, a number of points have been found in the course of the evaluation procedures that might be worth improving. We intend to address them all, among them:

1. Improving the readability of text: a clearer representation of mathematical formulas, adding some legends, and making existing legends more easily readable;
2. Improving the mouse handling (enabling the user to select graphical objects more easily);
3. Making it possible to define temporal reference points freely (i.e., removing the limitation of their being defined at the start of the treatment plan execution). This might be useful for investigating effects within long-term therapies (e.g., asthma therapy).

Conducting Further Evaluation

We have undertaken an number of steps to verify the quality and usefulness of CareCruiser (see Chapter 7). Still, it is advisable to subject CareCruiser to further evaluation, for instance:

1. Testing CareCruiser with a statistically significant number of users;
2. Testing CareCruiser with a statistically significant number of patient data;
3. Conducting long-term studies in the intended user environment (e.g., using it continually over an extended period of time in medical institutions).

Extending the Graphical Representation to Patient Cohorts

CareCruiser, for the time being, provides a possibility to make the user compare data derived from a limited number of patients only; 3-4 patients with expanded patient facets (see Figure 5.14) or about 15 patients with collapsed patient facets (see Figure 5.15). For the following reasons it would be advisable to overcome this limitation:

1. In the representation of a much greater number of patients, sub-groups of different types of patient behavior (with regard to the patients' reactions to the same treatment procedure) would become visible;
2. The greater the number of patients, the greater the significant relevance of their data;
3. In the context of what has been said above, this would enable the user to compare alternative therapies and their relative merits.

Investigating Possible Applications of Our Information Visualization Methods to Other Fields

As we have sketched out in Chapter 9, we believe our interactive InfoVis methods to be applicable with advantage to other fields as well. In this context we have been using simulated data only. Therefore, it is necessary to

1. Gather real data in other fields, such as approval ratings and newspaper events, sales figures and commercial actions, as well as stock market figures and events that influence the market;
2. Applying our visual methods to these fields;
3. Evaluating their usefulness in this context.

Probing into Event-Dependent Subjective Time Applications

Usually, time is measured in terms of hours, minutes, etc. However, there is a second way of measuring time: time units can be adjusted to events that take place. In other words, one unit of time might be the time that elapses between two successive events relating to one another. For instance, in biology the time between two relevant events, such as development stages, may differ in length, but they may – within these different lengths – exhibit similar features. In order to examine these features, it would make sense to arrange these development stages for closer comparison. Our interactive InfoVis methods as used in CareCruiser could be readily applied to align the starting events below each other and to adjust the individual time axes (i.e., lengthening or shortening the time axes) in a way that the follow-up events are also aligned below each other. This seems to be a very interesting research direction which we would like to probe into.

Our research about the visualization of effects of specific actions - especially the visualization and the interactive exploration of complex, time-oriented data - was quite a challenge. Meeting it has yielded some highly interesting and satisfactory results. Of particular interest was

to me the great power of Information Visualization to discover hitherto hidden patterns in the data. Our research pointed us to various further research areas (as mentioned above) that I find fascinating and well worth dealing with through future research activities. It sparked my interest to probe deeper into these exciting directions.

Part V

Appendix

User Study Protocol

Preamble

This protocol serves a two-fold purpose: (a) clearly stating all design decision for the user study that is to be conducted for the CareCruiser prototype, and (b) outlining the field procedures to carry the study through. Thus, it is used as a guide in conducting the user study. Data will be collected during and after the interaction phase where a medical expert uses the CareCruiser prototype to investigate a given dataset about the ventilation of new born children and to formulate hypothesis about treatment effects. Two observers will be present to document (taking notes) (1) general observations and (2) the findings of the think aloud method. Subsequent to this interaction phase, an interview will be conducted with the medical expert. Results of this user study will be published in a scientific paper. This protocol provides a brief overview of the research project followed by a description of the procedure for conducting the case. Moreover, we outline the research instruments we use to collect data and the procedures of analyzing the data.

General

Due to the use of modern information technologies in clinical care, the amount of information grows constantly. Thus, the application of Information Visualization (InfoVis) techniques to deal with the available information in a comprehensible way is imperative. By offering interactive visual representations, InfoVis aims at providing intelligibility and optimal use of the enormous amounts of complex and often interrelated data. Card et al. constituted InfoVis as ‘The use of computer-supported, interactive, visual representations of abstract data to amplify cognition’ [13]. In other words, ‘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’ [4]. Furthermore, it allows the user to explore data at different stages of complexity to avoid overload, e.g., providing means for both a compact overview and a thorough exploration

of the ‘information space’. InfoVis 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., patients’ data, or treatment processes, the visualization of computerized clinical practice guidelines is a challenge for InfoVis researchers. From the plan management point of view clinical guidelines can be seen as (time-oriented) plans or activities [57]. In contrast to this, patients’ data is composed of heterogeneous information spaces, i.e., including quantitative and qualitative data, high-frequency and low-frequency data, etc. Conceptually patients’ data can be seen as multidimensional information spaces. For both concepts a variety of InfoVis approaches have been developed in order to meet specific requirements [4]. For the visualization of a patient’s condition in combination with the assigned clinical guideline, though, only very few contributions [1, 5, 10, 86] exist.

‘Usually, the visual representations of clinical guidelines are simply combined with text-based explanations of the course of the patient’ [4]. However, visualizing this connection would be of great benefit, since changes of the patient’s condition in accordance with the execution of the therapy plan is of major interest to physicians. Still there is no appropriate InfoVis method to illustrate these dependencies. Thus, ‘...visualizing patients’ data in connection with clinical guidelines is still an open, but challenging issue’ [4].

The CareCruiser prototype is aimed at filling this gap. It is to visually clarify (1) in which way the application of specific therapy plans as well as the application of selected clinical actions cause changes in the health condition of a patient and (2) how well these treatment plans succeed in reaching the intended effects. The overall achievement of the project is intended to be the development of a method supporting the visualization of effects of processes on multidimensional data, with special attention to the execution of clinical guidelines.

1. The development of a InfoVis technique specifically suited to support the characteristics of the process of clinical practice guidelines in connection with a patient’s condition.
2. The derivation of methods generally applicable to visualize processes in combination with affected data.
3. The design of a user interface providing appropriate information and interactive features, in order to better support the different tasks emerging in the context of clinical guideline application.
4. The evaluation of strong points and shortcomings of the chosen InfoVis methods, the interactive techniques, and the design of the user interface.

Objective

The user study is aimed at answering the following research questions: Main question:

- Does the CareCruiser visualization support a detailed analysis of – and adequate reaction to – the interdependencies of a patient’s health condition and the applied treatment (plan)?
- Sub Questions:

- Does the CareCruiser prototype adequately visualize a patient's condition and its change over time in relation to the applied treatment plans?
- Does the CareCruiser prototype visualize adequate parameters with adequate degree of detail/abstraction of data concerning a patient's condition and the applied treatment plan(s)?
- Does the CareCruiser visualization support an effective investigation of effects of applied treatments and clinical actions?

Procedure

The CareCruiser prototype is used to visualize a treatment plan to ventilate newborn children. We provide patient data of four different patients which have been threatened with this treatment plan. Two physicians will use CareCruiser to investigate the patient data in combination with the applied treatment plan and interactively manipulate the presentation to explore effects of the applied treatment plan and effects of clinical actions. The interaction with the CareCruiser prototype will take about 2 hours and will be carried out by one medical expert at a time while two observers are documenting the findings.

Research instruments

Since the use of more than one data source – a technique known as triangulation – is highly recommended by many researchers [59,61,99] as a mechanism for increasing both the reliability and validity of qualitative research [16], we will collect data from different sources:

- Two observers will take notes while the medical expert is interacting with the prototype. One will note general observations, the other one will note the findings of the think aloud method.
- Subsequent to the interaction phase, one observer will conduct a semi-structured interview (a mix of open and closed questions) with the medical expert. The interview questions will be based on the research questions. The interview will be audio-recorded to make sure to record all the details which may get lost when taking notes only. After the interview, it will be transcribed into text. The transcripts will be reviewed by the interview subject to make sure he agrees with the interpretation of what was said before it is analyzed.

Transcribed Interviews

B.1 Interview I

1. Wo sehen Sie die Stärken von CareCruiser?

Dass man wichtige Parameter kontinuierlich überwachen kann, was notwendig ist. Ein typisches Beispiel in meinem Bereich wären die kindlichen Herztöne über die Zeit aufgetragen. Das ist eine Sache, die man einfach überwachen muss und je mehr man dabei unterstützt wird, desto besser. Was hier auch sehr schön ist, ist, dass man neben den Veränderungen der Patientenparameter auch alle Maßnahmen, die gesetzt wurden, in der Zeitachse sehen kann.

Insbesondere – wenn ich jetzt näher darüber nachdenke – sieht man natürlich auch die Effekte von Maßnahmen, die man vorher schon gesetzt hat, d.h. man kann sich auch anschauen, was man vorher schon ausprobiert hat und was schon erfolgreich war. Die Visualisierung der Veränderungen ist glaube ich auch sehr wertvoll wenn man z.B. ein Kind länger betreut oder das Kind an jemand anderen übergibt, dass man in ähnlichen Situationen nachschauen kann, wie das Kind auf Veränderungen der Einstellungen reagiert hat. Damit kann man dann Aktionen wählen, die sich schon bewährt haben, oder man tut etwas nicht noch einmal, was sich nicht bewährt hat. Da könnt ich mir vorstellen, dass das auch sinnvoll ist – gerade wenn man eine längere Zeitachse hat.

2. Wo sehen Sie Verbesserungspotential?

Ich tu mir ein bisschen schwer hier die mathematischen Formeln/Regeln zu verstehen. Z.B. die Berechnung der Frequenz ist nicht leicht nachzuvollziehen – wird die Frequenz erhöht oder verringert? Es ist von einem Arzt zu viel verlangt, dass er weiß, wenn es unter dem Bruchstrich mehr wird, die Zahl insgesamt kleiner wird und umgekehrt. Man könnte auch den Bruchstrich übersehen – vor allem wenn das wie hier [Anmerkung: im Prototypen im 'Logical View', siehe Grafik 5.4] in so einer Wurst steht. Die Darstellung

des Wissens über diese Regeln [Anmerkung: mathematische Formeln im ‘*Logical View*’] müsste glaub ich noch ein bisschen mundgerechter sein.

An und für sich gefällt mir das sehr gut mit diesem Baum hier, nur da gibt es keine Beschriftungen – nur wenn man mit der Maus auf den Knoten fährt. Also ich könnte mir vorstellen um meinen gesamten Handlungsspielraum - also die ganzen Regeln die es gibt – besser überblicken zu können, wäre Beschriftung hilfreich. Ich meine es ist ok, wenn man sich ein Detail herausklickt, dass man das dann im oberen Fenster die Details sehen kann, aber irgendwie hätte ich ein gutes Gefühl wenn ich den ganzen Handlungsspielraum auch besser überblicken könnte – das würde mir gefallen.

Am ‘Timeview’ gibt es eigentlich nicht viel auszusetzen. Mir würde wahrscheinlich die Funktion ‘Absolute Distance’ schon völlig genügen, aber die Funktionen ‘Slope’ und ‘Relative Distance’ sind auch nett zu haben. An diesem View gibt es eigentlich nicht viel zu verbessern – wüsst ich jetzt nicht was ich da besser machen könnte. Eventuell könnte man den Zielbereich des Parameters – wo man hin will – noch deutlicher hervorheben.

3. Glauben Sie, dass CareCruiser helfen kann die Entwicklung des Patienten über einen längeren Zeitraum schneller erfassen zu können?

Ja, absolut. Gerade wenn man einen längeren Zeitraum betrachtet sieht man die einzelnen Aktionen und was sie bewirken. Ich wüsste nicht wie man das sonst in irgendeiner Art und Weise ablesen könnte. Angenommen man hätte diese Visualisierung nicht, dann könnte man eigentlich gar nichts sagen über den Therapieverlauf, weil die Alternative wäre ja, dass man das textlich aufschreibt – getätigte Aktionen und eine Beschreibung der Veränderung des Patientenparameters – daran kann man kaum Effekte ablesen. Ich glaube diese Visualisierung ist überhaupt die einzige Methode diese Effekte von den gesetzten klinischen Aktionen überhaupt sichtbar zu machen.

4. Glauben Sie, dass CareCruiser helfen kann die Entwicklung des Patienten im direkten Zusammenhang mit der angewandten Therapie bzw. den klinischen Aktionen leichter erfassen zu können?

Man sieht es schon – was ein bisschen schwierig ist, ist wenn mehrere Aktionen gleichzeitig gesetzt werden, nachzuvollziehen woher jetzt ein spezieller Effekt kommt. Und manchmal findet man dann eben paradoxe Sachverhalte, z.B.: dass man eigentlich erwarten würde, dass jetzt ein gewisser Patientenparameter nach einer klinischen Aktion herunter geht, der aber zuerst hinauf geht – da findet man dann auch ganz unerwartete Dinge. Es ist dann im Nachhinein interessant was hier passiert ist – ich kann mir das jetzt nicht erklären. Da muss man sich dann noch andere Parameter anschauen – und eventuell klärt sich dann die Situation. Also man entdeckt sozusagen auch ganz paradoxe Reaktionen – vielleicht kann man auch daraus lernen. Da muss man dann die Situation noch genauer untersuchen und vielleicht findet man irgendeine Ursache dafür. Andere Patientenreaktionen nach gewissen Aktionen sind wiederum sehr logisch und man sieht sehr schön, dass eine gewisse klinische Aktion den erwarteten Erfolg hatte – ich würde dann, wenn ich einen Patienten übernehmen würde (wenn Schichtwechsel ist) und ich

sehe, dass gewisse klinische Aktionen sehr schön gewirkt haben, dann würde ich sehr wohl die gleichen Aktionen, die vorher schon erfolgreich waren, in einer ähnlichen Situation wiederholen. Das würde mir gefallen.

5. Glauben Sie, dass CareCruiser helfen kann klinische Situationen schneller einzuschätzen?

Schnell muss man bei kritischen Werten reagieren, weil dann muss man was tun. Das hätte man aber durch normale Alarmfunktionen auch. Da sehe ich jetzt nicht unbedingt den Vorteil – eher so, dass man sich Situationen die nicht so akut sind anschauen kann und da Adjustierungen vornehmen kann – Feinarbeiten. Dort sehe ich eher den Vorteil; in der Optimierung.

6. Glauben Sie, dass CareCruiser helfen kann den tendenziellen Verlauf von Therapien (positive/negative Wirkung) leichter erfassen zu können?

Oja, unbedingt. Mir ist jetzt auch grad eingefallen: z.B.: könnte man einem Kind auch ein Surfactant geben, dann würde man merken, dass sich durch diese Surfactant-Gabe sehr viele Dinge ändern. Dass z.B.: der Beatmungsdruck und der FiO₂ Bedarf gesenkt werden kann. Also, da würde man sehr schön sehen, dass wenn man da jetzt irgendetwas anderes tut, wie sich das auf die ganze Beatmungssituation auswirkt. Das wäre sehr instruktiv; das würde ich sehr lehrreich finden. Lehrreich nämlich einerseits für diesen einen Patienten, dass man weiß wie wirkt sich die Surfactant-Gabe auf die gesamte Beatmungssituation aus – andererseits kann man das auch sehr gut für die Lehre den Studenten demonstrieren wie sich so eine Surfactant-Gabe auswirkt. Auch die Entwicklung über eine längere Zeit: dass man sieht, nach der Surfactant-Gabe geht es dem Kind besser und über die Zeit geht es ihm dann wieder schlechter, dann müsste man nochmal ein Surfactant geben – dass man auch solche Dinge beurteilen kann.

7. Glauben Sie, dass CareCruiser helfen kann einzuschätzen, ob eine gewisse Therapie mehrheitlich den gewünschten Effekt hat (beim Betrachten einer größeren Patientengruppe)?

Man könnte Patienten miteinander vergleichen und schauen ob eine gewisse klinische Aktion bei allen die gleiche Wirkung hat oder nicht. Oder man könnte auch unterschiedliche Therapien (z.B.: unterschiedliche Beatmungsformen/Beatmungseinstellungen) vergleichen. Für die Lehre ist sowas natürlich sehr wertvoll, aber auch angenommen man macht eine Visite, dass man sich einfach mal die Beatmungssituation der Kinder anschaut und das einfach mal bespricht; die Kinder vergleichend bespricht.

8. Ist CareCruiser geeignet alle nötigen Informationen um den Patientenzustand einschätzen zu können in adäquater Form darzustellen?

Mir würde jetzt nichts abgehen. Was ich gerne hätte wäre, dass man auch andere Therapien die man gleichzeitig gemacht hat sehen kann [Anmerkung: das ist ohnehin in unserem Prototypen möglich]. Bei den klinischen Aktionen die nicht unmittelbar zur Ther-

apie gehören ist auf den ersten Blick nicht so ganz klar worum es geht – da wäre es gut nett wenn man lesen könnte was das für Aktionen sind (Labels). Etwas Ähnliches gibt es bei uns bei den Pathogrammen (mit Labels).

9. Ist CareCruiser geeignet alle nötigen Informationen der angewandten Therapie in adäquater Form darzustellen?

Also das ist extrem instruktiv (Parametereinstellungen / Parameterveränderungen). Das einzige was ein bisschen schwer ist, ist die Skalierung abzulesen – die Schrift ist etwas zu klein.

10. Glauben Sie, dass CareCruiser helfen kann nicht bekannte/ nicht offensichtliche Reaktionen auf klinische Aktionen zu erkennen?

Ja, absolut. Manchmal passiert aber nicht das Erwartete, das ist dann halt paradox.

11. Glauben Sie, dass CareCruiser helfen kann herauszufinden, ob gewisse Reaktionen bei einer Patientengruppe häufig bzw. weniger häufig auftreten?

Könnt ich mir schon vorstellen, dass man Reaktionen beobachtet. Z.B. bei Personengruppen mit unterschiedlichen Krankheiten. Was glaub ich auch besonders interessant ist, wenn ein Kind z.B. relativ stabil ist bei der Beatmung und dann verschlechtert es sich weil es krank wird: dass man das dann auch aus diesen Beatmungseinstellungen etwas über den Gesamtzustandes des Kindes erfährt. Z.B. dass es einfach kränker ist als es vorher war, oder auch gesünder. Die Kinder können z.B. relativ leicht eine Infektion bekommen – das merkt man dann oft an der Beatmung relativ früh, dass sie einfach schlechter zu beatmen sind (von den Beatmungseinstellungen). So eine Verschlechterung würde man bei einem Kind auch gut sehen im Verlauf.

12. Glauben Sie, dass CareCruiser helfen kann individuelle Reaktionen eines Patienten auf bestimmte klinische Aktionen erkennen zu können und aufgrund dieses Wissens die angewandte Therapie besser auf den einzelnen Patienten abstimmen zu können?

Oja, klar. Wenn bei einem Patienten die Therapie nicht so funktioniert wie bei den anderen Patienten muss man überlegen warum und dann die Therapie modifizieren eventuell noch andere Dinge zusätzlich machen oder eben eine andere Therapie anwenden.

13. Glauben Sie, dass CareCruiser helfen kann, zu überprüfen, ob der gewünschte Effekt nach einer klinischen Aktion mehrheitlich eintritt?

Ja, natürlich, wenn man sich mehrere Patienten anschaut.

B.2 Interview II

1. Wo sehen Sie die Stärken von CareCruiser?

Die Stärken sehe ich darin, dass die Visualisierung übersichtlich ist, dass sie handlich ist und auf wichtige Parameter begrenzt ist. Das ist eine sehr praktikable Lösung Daten zu analysieren mit interessanten Features. Also ich glaube, dass das eine sehr brauchbare Geschichte ist – intuitiv, wunderschön.

2. Wo sehen Sie Verbesserungspotential?

Vielleicht, dass die Veränderung zum Initialwert angezeigt wird, der aber in manchen Fällen nach einer halben Stunde nicht mehr so relevant ist. Da empfinde ich die ‘Total Distance’ Funktion für wichtiger – zumindest für die Beatmung. Für Langzeittherapien z.B. Asthma ist das aber schon sinnvoll. Bei der Beatmung sind unmittelbare Veränderungen wichtiger. Eine gute Möglichkeit wäre es wenn man den Referenzwert flexibel einstellen könnte.

Ich hätte vielleicht den Hierarchischen View oben und den Logischen View unten angeordnet, weil man im Hierarchischen View den Überblick sieht und im Logischen View die Details und ich ja von oben nach unten lese – das wäre für mich intuitiver. Aber das ist ein minimales Detail.

3. Glauben Sie, dass CareCruiser helfen kann die Entwicklung des Patienten über einen längeren Zeitraum schneller erfassen zu können?

Die Sinnhaftigkeit sich lange Zeitverläufe anzuschauen hängt vom Anwendungsgebiet ab. Bei der Beatmung geht es vielmehr um unmittelbare Reaktionen auf klinische Aktionen, da macht das Erfassen eines längeren Zeitraums weniger Sinn. Wenn man sich jetzt aber z.B. den Effekt einer Krebstherapie auf irgendwelche Zytostatika anschaut, dann ist das natürlich schon eine Möglichkeit sich auch lange Zeitverläufe anzuschauen. Da hat man ein bestimmtes Therapieschema und schaut wie entwickeln sich die Zahl der Tumorzellen oder wie entwickelt sich ein HIV Parameter auf dieses Therapieschema oder weiß der Kuckuck was - das wäre glaub ich auch ein Möglichkeit sich einen längeren Zeitraum anzuschauen. Es hängt glaub ich auch mehr von der Zahl der Aktionen ab – je weniger Aktionen desto längere Zeiträume kann man gut überblicken. Bei häufigen Aktionen und sehr raschen Änderungen dann ist die Visualisierung glaub ich eher eine kurzfristige Analyse geeignet.

Für einen groben Überblick könnte man natürlich mit Hilfe der Absolut Distance Funktion ablesen, dass der Patient z.B. am Anfang der Therapie kritisch und dann gegen Ende der Therapie eigentlich immer im Zielbereich war. Aber für die genauere Analyse hängt es von der Zahl der Aktionen ab wie lange Zeiträume man betrachtet, denn wenn man bei sehr vielen Aktionen über einen längeren Zeitraum jede Einzelne beurteilt, dann ermüdet man wahrscheinlich. Also, es macht Sinn um einen Überblick zu bekommen, aber um einzelne Aktionen zu beurteilen wäre eher die Analyse von kürzeren Zeiträumen sinnvoll.

4. **Glauben Sie, dass CareCruiser helfen kann die Entwicklung des Patienten im direkten Zusammenhang mit der angewandten Therapie bzw. den klinischen Aktionen leichter erfassen zu können?**

Ich glaub schon, ja. Allerdings weiß ich nicht ob man einen Erkenntnisgewinn hat, wenn man von der Materie gar nichts versteht [Anmerkung: ist für die Verwendung von Medizinern vorgesehen]. Aber ich glaube schon, dass man mit dieser Visualisierung aus einem Haufen Daten Regeln auch intuitiv findet.

5. **Glauben Sie, dass CareCruiser helfen kann klinische Situationen schneller einzuschätzen?**

Ich glaube, dass man die klinische Situation gut demonstrieren kann. Dass man z.B., seinen Schülern die Situation erklären kann - z.B. wie viel Zeit der Patient im nicht-optimalen Bereich verbracht hat, ob bei der Therapie geschlampt wurde oder ob auf gewisse Parameterwerte nicht richtig reagiert wurde - das sieht man schon alles. Das sind hier doch 10 Stunden, also ein ganzer Nachtdienst. Da kann man sehen, ob der Patient über die ganze Zeit optimal versorgt war oder nicht.

6. **Glauben Sie, dass CareCruiser helfen kann den tendenziellen Verlauf von Therapien (positive/negative Wirkung) leichter erfassen zu können?**

Sicher, ja. Bei der Beatmung ist zwar der Initialwert nicht so sehr relevant – das wäre ein flexibel wählbarer Referenzpunkt sinnvoller – aber bei anderen Therapien ist das sinnvoll. Und man kann sicher auch zeigen, ob der gewünschte Therapieerfolg eingetreten ist oder nicht. Manche Patienten bleiben instabil, aber natürlich sollte sich der Patient gegen Ende der Therapie stabilisiert haben.

7. **Glauben Sie, dass CareCruiser helfen kann einzuschätzen, ob eine gewisse Therapie mehrheitlich den gewünschten Effekt hat (beim Betrachten einer größeren Patientengruppe)?**

Da kann man sich dann im Überblick anschauen, ob eine Therapie den gewünschten Effekt hat. Wie oft ist die Therapie gut gegangen; wie oft ist der Patientenparameter gefallen nach einer bestimmten Aktion. Ich könnte mir vorstellen, dass man Rückschlüsse auf den Therapieeffekt bei einer bestimmten Patientengruppe ziehen kann. Bei der Beatmung ist das zwar nicht trivial aber bei anderen Dingen sicher. Man könnte z.B. sehen, dass die Gabe von Kortison bei so und so vielen Patienten eine Verbesserung gebracht hat und bei so und so vielen hat es eine Verschlechterung von dem und dem Parameter gebracht.

8. **Ist CareCruiser geeignet alle nötigen Informationen um den Patientenzustand einschätzen zu können in adäquater Form darzustellen?**

Ja, die Patientendaten reichen vollkommen aus.

9. **Ist CareCruiser geeignet alle nötigen Informationen der angewandten Therapie in adäquater Form darzustellen?**

Die Informationen reichen auch aus. Es ist ein bisschen schwierig mit der Maus genau auf die Raute zu fahren damit man die Detail-Fenster bekommt, aber im Prinzip ist es ok – ich kann es mir nicht anders vorstellen. Nein, das ist wunderbar.

10. Glauben Sie, dass CareCruiser helfen kann nicht bekannte/ nicht offensichtliche Reaktionen auf klinische Aktionen zu erkennen?

Im Prinzip ja. Wenn ich sehe ich erhöhe da irgendwas und es tut sich gar nichts, sollte es mich nachdenken lassen, ob da irgendwas nicht richtig gemacht worden ist. Ich würde dann andere Parameter untersuchen und nach der Ursache suchen – hab ich auch gemacht.

11. Glauben Sie, dass CareCruiser helfen kann herauszufinden, ob gewisse Reaktionen bei einer Patientengruppe häufig bzw. weniger häufig auftreten?

Dazu ist die Beatmung vielleicht zu komplex, aber bei anderen Therapien könnt ich mir das vorstellen. Z.B. bei der Chemotherapie oder bei der Psychotherapie oder medikamentöse Therapie. Dass ich mir das anschau auf den Depressionsscore zum Beispiel – ich gebe ein Antidepressivum und sehe nach acht Wochen bessert es sich – da könnt ich mir schon vorstellen dass man erkennen kann ob eine gewisse Patientengruppe anders reagiert als andere. Es stellt sich halt die Frage wie viele Patienten kann man sich parallel anschauen – wahrscheinlich hängt es von der Anzahl der Parameter ab, aber ich glaub eher, dass es für Gruppenvergleiche ganz gut geeignet ist. Aber im Prinzip kann ich mir das schon vorstellen. Und da würde es auch Sinn machen vom Initialwert – beim Depressionsscore zum Beispiel, wenn der bei 35 anfängt – wie geht es dann nach acht Wochen? Oder Psychose-Score wie ‘brief psychiatric rating scale’ – BPRS – du gibst ihm das Mittel und 5 Tage später – und dann hat er irgendein ‘Adverse Event’ oder so und dann geht er wieder hinauf... Kann ich mir schon vorstellen, ja.

Man könnte ja auch Mittelwerte von Gruppen vergleichen – wenn man die Gruppen vorher auswählt: das sind Responder, das sind Nicht-Responder, kann man auch Responder und Nicht-Responder vergleichen. Das würde ich auch so sehen. Das wäre sicher auch eine Anwendungsmöglichkeit, dass man Gruppenvergleiche macht.

12. Glauben Sie, dass CareCruiser helfen kann individuelle Reaktionen eines Patienten auf bestimmte klinische Aktionen erkennen zu können und aufgrund dieses Wissens die angewandte Therapie besser auf den einzelnen Patienten abstimmen zu können?

Sicher. Das hängt halt davon ab, ob man alle nötigen Parameter die man braucht um diesen Schluss zu ziehen zur Verfügung hat. Wenn ich z.B. bemerke, nach irgendeinem speziellen Mittel rasseln die Thrombozyten in den Keller oder der Patient geht mit den Transaminasen hinauf dann kann man dementsprechend die Therapie modifizieren. Aber wenn man die Transaminasen gar nicht erfasst dann geht’s natürlich nicht. Aber wenn alle nötigen Informationen vorhanden sind, dann kann ich mir schon vorstellen, dass man die Therapie mit diesem Wissen optimiert.

13. **Glauben Sie, dass CareCruiser helfen kann, zu überprüfen, ob der gewünschte Effekt nach einer klinischen Aktion mehrheitlich eintritt?**

Ich denke schon, ja.

Weitere Anmerkungen: Es ist wunderschön geworden – wirklich. Durchaus intuitiv bedienbar – sehr sehr schön. Auch sozusagen die Idee das man die Pläne darstellt und das man dann den Plan den man im einen View auswählt auch in den anderen Views sieht – sehr schön, super.

Bibliography

- [1] W. Aigner. Guideline overview tool (got), Technical Report Asgaard-TR-2002-4. Technical report, Vienna University of Technology, Institut of Software Technology & Interactive Systems, Vienna, Austria, 2001.
- [2] W. Aigner. Interactive visualization of time-oriented treatment plans and patient data. MSc thesis in computer science, Vienna University of Technology, 2003.
- [3] W. Aigner, P. Federico, T. Gschwandtner, S. Miksch, and A. Rind. Tackling the challenges of time-oriented visual analytics in healthcare. In *Proceedings of the Workshop on Visual Analytics in Healthcare (VAHC 2012) in conjunction with IEEE VisWeek 2012, Oct 14-19, Seattle, WA, USA*, page accepted, 2012.
- [4] W. Aigner, K. Kaiser, and S. Miksch. Visualization methods to support guideline-based care management. In A. ten Teije, S. Miksch, and P. J. F. Lucas, editors, *Computer-Based Medical Guidelines and Protocols: A Primer and Current Trends*, volume 139 of *Studies in Health Technology and Informatics*, pages 140–159. IOS Press, Amsterdam, The Netherlands, 2008.
- [5] W. Aigner and S. Miksch. CareVis: Integrated visualization of computerized protocols and temporal patient data. *Artificial Intelligence in Medicine (AIIM)*, 37(3):203–218, 2006.
- [6] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-Oriented Data*. Springer-Verlag, Berlin / Heidelberg, 1st edition, 2011.
- [7] W. Aigner, S. Miksch, B. Thurnher, and S. Biffl. PlanningLines: Novel glyphs for representing temporal uncertainties and their evaluation. In *Proceedings of the 9th International Conference on Information Visualisation (IV05), July 6-8, London, UK*, pages 457–463, Los Alamitos, CA, USA, 2005. IEEE Computer Society Press.
- [8] F. Anscombe. Graphs in statistical analysis. *The American Statistician*, 27(1):17–21, 1973.
- [9] R. Bade. Methoden zur Visualisierung von und Interaktion in zeitbasierten Patientendaten und Behandlungsplänen. MSc thesis in computer science, University of Magdeburg, 2002.
- [10] R. Bade, S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2004), April 24-29, Vienna, Austria*, pages 105–112, New York, NY, 2004. ACM Press.

- [11] M. Burch, F. Beck, and S. Diehl. Timeline trees: visualizing sequences of transactions in information hierarchies. In *Proceedings of the working conference on Advanced Visual Interfaces (AVI 2008), May 28-30, Napoli, Italy*, pages 75–82, New York, NY, 2008. ACM Press.
- [12] S. Card. Information visualization. In J. Jacko and A. Sears, editors, *The human-computer interaction handbook*, pages 509–543. Lawrence Erlbaum Associates Inc., Hillsdale, NJ, USA, 2007.
- [13] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, Burlington, MA, 1999.
- [14] S. Carpendale. Evaluating information visualizations. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization: Human-Centered Issues and Perspectives*, Lecture Notes in Computer Science, pages 19–45. Springer-Verlag, Berlin / Heidelberg, 2008.
- [15] J. Chambers, W. Cleveland, B. Kleiner, and P. Tukey. *Graphical Methods for Data Analysis*. Wadsworth Publishing Company, Belmont, CA, 1983.
- [16] P. Chau. On the use of construct reliability. *MIS Research: A Meta Analysis. Information and Management*, 35:217–227, 1999.
- [17] L. Chittaro, C. Combi, and G. Trapasso. Data mining on temporal data: A visual approach and its clinical application to hemodialysis. *Journal of Visual Languages and Computing*, 14(6):591–620, 2003.
- [18] P. Ciccarese, E. Caffib, L. Boiocchia, A. Halevy, S. Quaglinia, A. Kumar, and M. Stefanelli. The NewGuide Project: guidelines, information sharing and learning from exceptions. In *Proceedings of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), October 18-22, Protaras, Cyprus*, pages 163–167, Berlin / Heidelberg, 2003. Springer-Verlag.
- [19] P. Ciccarese, E. Caffib, L. Boiocchia, S. Quaglinia, and M. Stefanellia. A guideline management system. *Stud Health Technol Inform*, 107(Pt 1):28–32, 2004.
- [20] Clinical Practice Guideline on Major Depression in Childhood and Adolescence. http://www.guiasalud.es/egpc/traduccion/ingles/depression_infancia/completa. (Accessed: 2012-07-11).
- [21] CVASt – Laura Bassi Centre for Visual Analytics Science and Technology. <http://www.cvast.tuwien.ac.at/>. (Accessed: 2009-08-12).
- [22] P. A. de Clercq, J. A. Blom, H. H. M. Korsten, and A. Hasman. Approaches for creating computer-interpretable guidelines that facilitate decision support. *Artificial Intelligence in Medicine (AIIM)*, 31(1):1–27, 2004.
- [23] P. Eades. A heuristic for graph drawing. *Congressus Numerantium*, 42:149–160, 1984.

- [24] G. Eilenberger. Komplexität. Ein neues Paradigma der Naturwissenschaften. Cf. *Mannheimer Forum*, 89/90:71–136, 1999.
- [25] P. L. Elkin, M. Peleg, R. Lacson, E. Bernstam, S. W. Tu, A. Boxwala, R. Greenes, and E. H. Shortliffe. Toward standardization of electronic guideline representation. *MD Computing*, 17(6):39–44, 2000.
- [26] J.-D. Fekete, J. J. van Wijk, J. T. Stasko, and C. North. The value of information visualization. *Information Visualization*, 4950:1–18, 2008.
- [27] S. Few. *Show Me the Numbers: Designing Tables and Graphs to Enlighten*. Analytics Press, Burlingame, CA, 2004.
- [28] M. J. Field and K. N. Lohr. *Clinical Practice Guidelines: Directions for a New Program*. National Academies Press, Washington, D.C., 1990.
<http://www.nap.edu/books/0309043468/html/index.html> (Accessed: 2009-07-15).
- [29] Society for Medical Decision Making. Proposal for clinical algorithm standards. *Medical Decision Making*, 12(2):149–154, 1992.
- [30] C. Forsell and J. Johansson. An heuristic set for evaluation in information visualization. In *Proceedings of the international conference on Advanced Visual Interfaces (AVI 2010)*, May 25-29, Roma, Italy, pages 199–206, New York, NY, 2010. ACM Press.
- [31] J. H. Gennari, M. A. Musen, R. W. Ferguson, W. E. Grosso, M. Crubézy, H. Eriksson, N. F. Noy, and S. W. Tu. The evolution of Protégé: An environment for knowledge-based systems development. *International Journal of Human-Computer Studies*, 58:89–123, 2003.
- [32] J. Gleick. *Chaos: Making a New Science*. Viking Penguin, New York, NY, 1987.
- [33] K. Gödel. Zum intuitionistischen Aussagenkalkül (On the intuitionistic propositional calculus). *Anzeiger der Akademie der Wissenschaften in Wien*, 69:65–66, 1932.
- [34] H. H. Goldstine and J. von Neumann. *Planning and coding problems for an electronic computing instrument, part II*, volume 1. U.S. Ordinance Department, 1947. (Reprinted in von Neumann, J. 1963. *Collected Works*, volume 5. New York: McMillian, pp. 80-151).
- [35] T. Gschwandtner. Interactive visualization to optimize treatment choices for individual patients. In *IEEE VisWeek 2009 Doctoral Colloquium, October 10, Atlantic City, New Jersey, USA*, 2009.
- [36] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. CareCruiser: Exploring and visualizing plans, events, and effects interactively. In *Proceedings of the IEEE Pacific Visualization Symposium (PacificVis 2011)*, March 1-4, Hong Kong, P. R. of China, pages 43–50, Los Alamitos, CA, 2011. IEEE Computer Society Press.

- [37] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. Design and evaluation of an interactive visualization of therapy plans and patient data. In *Proceedings of the 25th BCS Conference on Human Computer Interaction (HCI2011), July 4-8, Newcastle upon Tyne, UK*, pages 421–428, London, UK, 2011. BCS.
- [38] T. Gschwandtner, J. Gärtner, W. Aigner, and S. Miksch. A taxonomy of dirty time-oriented data. In G. Quirchmayr, J. Basl, I. You, L. Xu, and E. Weippl, editors, *Multidisciplinary Research and Practice for Information Systems (Proceedings of the CD-ARES 2012), Aug 20-24, Prague, Czech Republic*, volume 7465 of *Lecture Notes in Computer Science*, pages 58–72. Springer-Verlag, Berlin / Heidelberg, 2012.
- [39] T. Gschwandtner, K. Kaiser, and S. Miksch. Information requisition is the core of guideline-based medical care: Which information is needed for whom? *Journal of Evaluation in Clinical Practice (JECP)*, 17:713–721, 2011.
- [40] D. C. Hadorn. Use of algorithms in clinical practice guideline development: Methodology perspectives. *Clinical Practice Guideline Development: Methodology Perspectives*, 9(95):93–104, 1995.
- [41] M. Harrower and C. Brewer. Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.
- [42] J. Heer, M. Bostock, and V. Ogievetsky. A tour through the visualization zoo <http://hci.stanford.edu/jheer/files/zoo/>. (Accessed: 2011-08-05).
- [43] K. Hinum, S. Miksch, W. Aigner, S. Ohmann, C. Popow, M. Pohl, and M. Rester. Gravi++: Interactive information visualization of highly structured temporal data. *Journal of Universal Computer Science, Special Issue on Visual Data Mining*, 11:1792–1805, 2005.
- [44] H. Hochheiser and B. Shneiderman. Interactive exploration of time series data. In K. Jantke and A. Shinohara, editors, *Discovery Science*, volume 2226 of *Lecture Notes in Computer Science*, pages 441–446. Springer-Verlag, Berlin / Heidelberg, 2001.
- [45] A. Holzinger. Usability engineering methods for software developers. *Communications of the ACM*, 48(1):71–74, 2005.
- [46] W. Horn, C. Popow, and L. Unterasinger. Metaphor graphics to visualize ICU data over time. In *Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-98), Workshop Notes of the ECAI-98 Workshop, August 25, Brighton, UK*, pages 76–81, 1998.
- [47] InferMed. <http://www.infermed.com/index.php/arezzo>. (Accessed: 2012-07-11).
- [48] D. A. Keim, J. Schneidewind, and M. Sips. CircleView a new approach for visualizing timerelated multidimensional data sets. In M.F. Costabile, editor, *Proceedings of the working conference on Advanced visual interfaces (AVI '04), May 25-28, Gallipoli, Italy*, pages 179–182, New York, NY, 2004. ACM Press.

- [49] KeysCorner. <http://www.keyscorner.com/archives/2005/02/13/recipe-diagramming/>. (Accessed: 2009-01-06).
- [50] R. Kincaid and H. Lam. Line graph explorer: Scalable display of line graphs using focus+context. In *Proceedings of the working conference on Advanced Visual Interfaces (AVI 2006), May 23-26, Venezia, Italy*, pages 404–411, New York, NY, 2006. ACM Press.
- [51] D. Klimov, Y. Shahar, and M. Taieb-Maimon. Intelligent visualization and exploration of time-oriented data of multiple patients. *Artificial Intelligence in Medicine (AIIM)*, 49(1):11–31, 2010.
- [52] R. Kosara and S. Miksch. Metaphors of movement: A visualization and user interface for time-oriented, skeletal plans. *Artificial Intelligence in Medicine (AIIM), Special Issue: Information Visualization in Medicine*, 22(2):111–131, 2001.
- [53] J. Lamping, R. Rao, and P. Pirolli. A focus context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 1995), May 7-11, Denver, Colorado, USA*, pages 401–408, New York, NY, 1995. ACM Press.
- [54] C. Lewis and J. Rieman. *Task-Centered User Interface Design - A Practical Introduction*. C. Lewis and J. Rieman, University of Colorado, Boulder, CO, 1993. <http://hcibib.org/tcuid/> (Accessed: 2011-12-22).
- [55] X. Lin. Map displays for information retrieval. *Journal of the American Society for Information Science*, 48(1):40–54, 1997.
- [56] P. McLachlan, T. Munzner, E. Koutsofios, and S. North. LiveRAC - interactive visual exploration of system management time-series data. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2008), April 5-10, Florence, Italy*, pages 1483–1492, New York City, NY, USA, 2008. ACM Press.
- [57] S. Miksch. Plan management in the medical domain. *AI Communications*, 12(4):209–235, 1999.
- [58] S. Miksch, Y. Shahar, W. Horn, C. Popow, F. Paky, and P. Johnson. Time-oriented skeletal plans: Support to design and execution. In *Fourth European Conference on Planning (ECP'97), September 24-26, Toulouse, France*, pages 299–311, Berlin / Heidelberg, 1997. Springer-Verlag.
- [59] M. B. Miles and M. Huberman. *Qualitative Data Analysis: An Expanded Sourcebook*. Sage Publications, Thousand Oaks, CA, 2nd edition, 1994.
- [60] I. Nassi and B. Shneiderman. Flowchart techniques for structured programming. *ACM SIGPLAN Notices*, 8(8):12–26, 1973.
- [61] W. L. Neuman. *Social Research Methods: Qualitative and Quantitative Approaches*. Allyn And Bacon, Boston, MA, 4th edition, 2000.

- [62] J. Nielsen. Heuristic evaluation. In J. Nielsen and R.L. Mack, editors, *Usability inspection methods*, number 38, pages 25–62. John Wiley & Sons, Inc., New York, NY, 1994.
- [63] R. Orsini. http://www.oreillynet.com/ruby/blog/2006/07/treemap_on_rails.html. (Accessed: 2009-01-09).
- [64] M. Peleg, S. W. Tu, J. Bury, P. Ciccarese, J. Fox, R. A. Greenes, R. Hall, P. D. Johnson, N. Jones, A. Kumar, S. Miksch, S. Quaglini, A. Seyfang, E. H. Shortliffe, and M. Stefanelli. Comparing computer-interpretable guideline models: A case-study approach. *Journal of the American Medical Informatics Association (JAMIA)*, 10(1):52–68, 2003.
- [65] K. Peters. <http://home.southernct.edu/~petersk1/csc400/csc400-flowchart.htm>. (Accessed: 2009-01-05).
- [66] C. A. Petri. Communication with automata, Technical Report RADC-TR-65-377. Technical report, Institut für Instrumentelle Mathematik, Bonn, Germany, 1966.
- [67] C. A. Petri and W. Reisig. http://www.scholarpedia.org/article/petri_net. (Accessed: 2011-08-05).
- [68] C. Plaisant, S. Lam, B. Shneiderman, M. Smith, D. Roseman, G. Marchand, M. Gillam, C. Feied, J. Handler, and H. Rappaport. Searching electronic health records for temporal patterns in patient histories: A case study with microsoft amalga, Technical Report HCIL-2008-13. Technical report, University of Maryland, College Park, Maryland, 2008.
- [69] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Lifelines: visualizing personal histories. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 1996), April 13-18, Vancouver, British Columbia, Canada*, pages 221–227, New York City, NY, USA, 1996. ACM Press.
- [70] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, and B. Shneiderman. Lifelines: Using visualization to enhance navigation and analysis of patient records. In *Proceedings of the AMIA Symposium (AMIA 1998), November 7-11, Orlando, Florida, USA*, pages 76–80, Philadelphia, PA, 1998. Hanley & Belfus, Inc.
- [71] S. Powsner and E. Tufte. Graphical summary of patient status. *The Lancet*, 344:386–389, 1994.
- [72] S. Quaglinia, M. Stefanelli, G. Lanzola, V. Caporusso, and S. Panzarasa. Flexible guideline-based patient careflow systems. *Artificial Intelligence in Medicine (AIIM)*, 22(1):65–80, 2001.
- [73] W. Randelshofer. <http://www.randelshofer.ch/treenviz/index.html>. (Accessed: 2009-01-09).
- [74] ReMINE – A high performances prediction, detection, and monitoring platform for patient safety risk management. <http://www.remine-project.eu/>. (Accessed: 2009-08-12).

- [75] G. Robertson, R. Fernandez, D. Fisher, B. Lee, and J. Stasko. Effectiveness of animation in trend visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1325–1332, 2008.
- [76] L. Rosenblum, R.A. Earnshaw, J. Encarnacao, H. Hagen, A. Kaufman, S. Klimenko, G. Nielson, F. Post, and D. Thalmann. *Scientific Visualization: Advances and Challenges*. Academic Press, Waltham, MA, 1994.
- [77] A. Seyfang, K. Kaiser, T. Gschwandtner, and S. Miksch. Visualizing complex process hierarchies during the modeling process. In *Proceedings of the 1st International Workshop on Theory and Applications of Process Visualization (TAProViz’12), September 3, Tallinn, Estonia*, page accepted, New York, NY, 2012. ACM Press.
- [78] A. Seyfang, R. Kosara, and S. Miksch. Asbru 7.3 reference manual, Technical Report Asgaard-TR-2002-1. Technical report, Vienna University of Technology, Institut of Software Technology & Interactive Systems, Vienna, Austria, 2002.
- [79] Y. Shahar and C. Cheng. Knowledge-based visualization of time-oriented clinical data. In *Proceedings of the AMIA Symposium (AMIA 1998), November 7-11, Orlando, Florida, USA*, pages 155–159, Philadelphia, PA, 1998. Hanley & Belfus, Inc.
- [80] Y. Shahar, D. Goren-Bar, M. Galperin, D. Boaz, and G. Tahan. Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artificial Intelligence in Medicine (AIIM)*, 38(2):115–135, 2006.
- [81] Y. Shahar, S. Miksch, and P. Johnson. The Asgaard Project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial Intelligence in Medicine (AIIM)*, 14(1):29–51, 1998.
- [82] Y. Shahar, O. Young, E. Shalom, A. Mayaffit, R. Moskovitch, A. Hessing, and M. Galperin. DECEL: A hybrid, multiple-ontology framework for specification and retrieval of clinical guidelines. In M. Dojat, E. Keravnou, and P. Barahona, editors, *Proceedings of the 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), October 18-22, Protaras, Cyprus*, volume 2780 of LNAI, pages 122–131, Berlin / Heidelberg, 2003. Springer-Verlag.
- [83] B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics (TOG)*, 11(1):92–99, 1992.
- [84] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages, September 3-6, Boulder, Colorado, USA*, pages 336–343, Los Alamitos, CA, 1996. IEEE Computer Society Press.
- [85] I. Spence and R. F. Garrison. A remarkable scatterplot. *The American Statistician*, 47(1):12–19, 1993.

- [86] R. Steele and J. Fox. Tallis PROforma Primer - introduction to PROforma language and software with worked examples. Technical report, Advanced Computation Laboratory, Cancer Research, London, UK, 2002.
- [87] D. Stotts. <http://www.cs.unc.edu/stotts/145/crc/state.html>. (Accessed: 2009-01-05).
- [88] P. Terenziani, S. Montani, A. Bottrighi, M. Torchio, G. Molino, and G. Correndo. The GLARE approach to clinical guidelines: Main features. *Studies in Health Technology and Informatics*, 101:162–166, 2004.
- [89] The Dutch Institute for Healthcare Improvement. <http://www.cbo.nl>. (Accessed: 2009-08-12).
- [90] The Euler Archive. <http://eulerarchive.maa.org/historica/euler-timeline.html>. (Accessed: 2011-08-05).
- [91] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Press, 2005.
- [92] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, 1983.
- [93] E. R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, CT, 1990.
- [94] S. Twaddle. Clinical practice guidelines. *Singapore Medical Journal*, 46(12):681–687, 2005.
- [95] T. D. Wang, C. Plaisant, A. Quinn, R. Stanchak, B. Shneiderman, and S. Murphy. Aligning temporal data by sentinel events: Discovering patterns in electronic health records. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2008), April 5-10, Florence, Italy*, pages 457–466, New York City, NY, USA, 2008. ACM Press.
- [96] Wikipedia. http://en.wikipedia.org/wiki/gantt_chart. (Accessed: 20011-12-22).
- [97] Wikipedia. http://en.wikipedia.org/wiki/state_diagram. (Accessed: 2011-08-05).
- [98] Wikipedia. Neonatal intensive care unit – http://en.wikipedia.org/wiki/neonatal_intensive_care_unit. (Accessed: 2012-07-11).
- [99] R. K. Yin. *Case study research. Design and methods*. Sage Publications, Thousand Oaks, CA, 2nd edition, 1994.