

CareCruiser: Exploring and Visualizing Plans, Events, and Effects Interactively

Theresia Gschwandtner* Wolfgang Aigner† Katharina Kaiser‡ Silvia Miksch§ Andreas Seyfang¶

Institute of Software Technology & Interactive Systems (ISIS)
Vienna University of Technology, Austria

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 a prototype that enhances known visualization methods to communicate the processes of treatment plan application together with their effects on a patient’s condition in an easily understandable way. Our prototype combines the advantages of enhanced visual recognition of patterns with traditional information of parameters’ development. 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.

Index Terms: H.5.m [Information Systems]: Information Interfaces And Presentation (e.g., HCI)—Miscellaneous; I.3.6 [Computing Methodologies]: Computer Graphics—Methodology and Techniques; J.3 [Computer Applications]: Life and Medical Sciences—Medical information systems

1 INTRODUCTION

Clinical practice guidelines (CPGs) [6] are text documents that provide recommendations for specific clinical situations. These recommendations are derived from best available scientific evidence and expert opinions. CPGs are modeled into computer-executable treatment plans containing precise recommendations of clinical actions tailored to specific clinical situations. Thus, they are powerful tools to support clinical decision-makers with the best available scientific evidence at the point of care.

When dealing with computer-executable treatment plans (e.g., applying them to a patient), the complex nature of these treatment plans calls for a plain and compact visualization of the underlying information. Dealing with logical sequences, hierarchical data, as well as time-oriented data, the visualization of clinical treatment plans relates to several specific fields of Information Visualization. Moreover CPGs are constantly improved to keep them up to date.

Improving these CPGs comprise several difficult tasks. Assessing the quality of the CPGs is a crucial precondition. In particular, there is a need to investigate the actual effects of the application of

different treatment plans and clinical actions to optimize the choice of treatment actions. In order to meet these needs, we have developed an interactive visualization (called CareCruiser) based on existing visualization techniques [3] for assessing effects of treatment plan application in combination with patient data interactively. Our visualization allows medical experts to more easily investigate the effects of previously applied treatment plans and clinical actions on the patient’s condition and thus it enables them to improve the quality of care by readjusting CPGs with respect to better information. In particular CareCruiser provides several features to support a step-wise interactive exploration of the patient’s condition and effects of applied treatment plans:

- Visualizing patient data in combination with applied treatment plans and clinical actions
- Aligning treatment plans and clinical actions for comparison
- Color-coded highlighting of interesting events of the patient’s parameter values’ development
- Filtering these highlighted events
- Providing focus and context information to support the detection of effects’ patterns

To start with, we give an outline of related work in Section 2. In Section 3 we describe the different types of information that demand special consideration. We continue with a detailed description of our visual encodings in Section 4. We outline the evaluation of our prototype and the extensions and improvements we intend to implement in Section 5. Finally, we sum up the main results of our work in Section 6.

2 RELATED WORK

Various approaches of visualizing treatment plans and patient data exist. However, visualizations of applied treatment plans in combination with patient data are found less often [2].

The first group of visualization approaches focuses on the visualization of treatment plans. Clinical Algorithm Maps [11], the Tallis toolset [25], Protégé [8], GUIDE [23], and Glare [26] offer a flowchart-like representation of clinical algorithms suited to visualize the logics and execution sequence of plans, but neglect the hierarchy of nested treatment plans and temporal constraints. AsbruView [18] is able to visualize all relevant plan characteristics (such as various constraints, the nesting of plans, and temporal aspects), but does not visualize patient data.

On the other hand, there is a variety of approaches visualizing patient data. The Graphical Summary of Patient Status [22] draws a comprehensive picture of a patient’s condition over time; LifeLines [21] visualizes qualitative and quantitative parameters of patient records, while LifeLines2 [27] and PatternFinder [20] are focused on filtering patient data and visualizing a set of special events to reveal interesting patterns. KNAVE II [24] represents different patient parameters and allows the user to explore them by means of

*e-mail: gschwandtner@ifs.tuwien.ac.at

†e-mail: aigner@ifs.tuwien.ac.at

‡e-mail: kaiser@ifs.tuwien.ac.at

§e-mail: miksch@ifs.tuwien.ac.at

¶e-mail: seyfang@ifs.tuwien.ac.at

multiple features, like different abstraction levels, absolute and relative time scales, different granularities of time, additional statistics, etc. VISITOR [17] is aimed at visualizing parameters of multiple patient records. VIE-VISU [15] is designed for the very special purpose of visualizing neonatal intensive care data. It uses glyphs to represent the patient's condition; unfortunately, the correct interpretation of these glyphs demands a huge learning effort. The 3D visualization technique of the Interactive Parallel Bar Charts approach [5] allows for exploring clinical parameters of hemodialysis sessions. This technique is quite space consuming; besides, problems of 3D representations (e.g., occlusions) are still not solved satisfactorily. Gravi++ [13] visualizes the change of the patient's condition over time by using animation and traces. However, none of these approaches is aimed at communicating characteristics of treatment plan application.

There are only a few tools visualizing applied treatment plans in combination with patient data. The Guideline Overview Tool [1] shows basic characteristics of treatment plan application in combination with patient data, Midgaard [4] visualizes detailed and compact patient data as well as the hierarchical structure of plans, complex characteristics (e.g., plan states and conditions), and execution sequences. CareVis [3] uses visualization techniques to represent relevant characteristics and temporal constraints of treatment plans in combination with patient data.

One sophisticated approach of visualizing large collections of timeseries data is LiveRAC [19]. It arranges a huge number of charts in a matrix and provides features such as interactive reordering 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.

The Line Graph Explorer [16] uses different colors, color saturation, and luminance to reduce the y-dimension of line charts to 1D ribbons and thus, to visualize large collections of line graph data at limited display space. We use a similar approach to code value ranges of multiple line charts. However, our color-codings take advantage of semantic information such as treatment plan goals and parameter progress since the start of the treatment. 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.

Still, 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.

3 PROBLEM ANALYSIS AND REQUIRED TYPES OF INFORMATION

In the following we give a use case and outline the required types of information.

Use Case An exemplary application of a treatment plan would be the ventilation of new born children. This treatment plan aims at stabilizing the patient's transcutaneous assessed oxygen saturation ($tcpSO_2$) 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 $tcpSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 mmHg.

Moreover, the FiO_2 is to be kept below 40% and the peak inspiratory pressure (PIP) below 60 cm H_2O . This treatment plan can be decomposed into sub-plans of which each may aim at achieving selected aspects of the overall intentions, for instance, the sub-plan 'Handle $tcpSO_2$ ' aims at stabilizing the patient's $tcpSO_2$ between 90% and 92% by adjusting the supply of FiO_2 .

Required Types of Information Our visualization aims at supporting the physician in monitoring the progress of one or more patients toward treatment plan intentions as well as identifying critical values at the first sight. Moreover, the visualization supports 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.

When visualizing the execution of such a treatment plan, there are two interconnected sets of data and information to be considered:

1. The progress of the treatment plan and applied clinical actions, and
2. The condition of the patient at any given point of treatment.

On the one hand, it is crucial to visualize the applied treatment plans (what therapy or medication has been applied, which clinical actions have been applied, etc). Besides the progress of the treatment, the internal structure of these plans is to be communicated by the visualization. Treatment plans usually consist of different sub-plans, which again may contain sub-plans, etc. These plans and sub-plans have to be executed in a strictly defined order, which should be evident as well. On the other hand, the course of the patient in connection with the applied treatment plans has to be represented in an intelligible way.

We have conducted a detailed analysis [10] of the visualization-relevant characteristics of these two interconnected sets of data, focusing on (1) the application of treatment plans and (2) the corresponding condition of the patient. This analysis has resulted in the following list of requirements:

1. Visualizing hierarchical data: the nesting of treatment plans and sub-plans
2. Visualizing temporal data: the execution sequence of treatment plans as well as the patient's condition and its evolution over time
3. Visualizing qualitative data: relevant characteristics of treatment plans, such as intentions of the plan, abort- and complete conditions etc.
4. Providing interactive means to allow for an active investigation of the course of the patient's condition and the detection of effects of applied treatments, and
5. Providing means to compare multiple patients.

In order to provide a tool that supports every aspect of treatment plan execution, we decided to extend the possibilities of CareVis which already covers point 1. to point 3. quite well (i.e., the nesting of treatment plans, the execution order, relevant plan characteristics, and temporal constraints). For this reason we have developed a conceptual extension of CareVis, namely CareCruiser, which, on the one hand, provides means to explore and assess success or failure of applied treatment plans and effects of clinical actions on a patient's condition (point 4.), and, on the other hand, allows for the investigation and comparison of two or more patients at the same time (point 5.).

4 DESIGN DECISIONS

Extending the possibilities of CareVis[3] resulted in adding value to an already very powerful tool, and thus, to provide a tool to support every single aspect important in the application of computer-executable treatment plans. On the other hand, it was a great challenge to find ways to smoothly integrate the visual encoding of additional information into the user interface of CareVis.

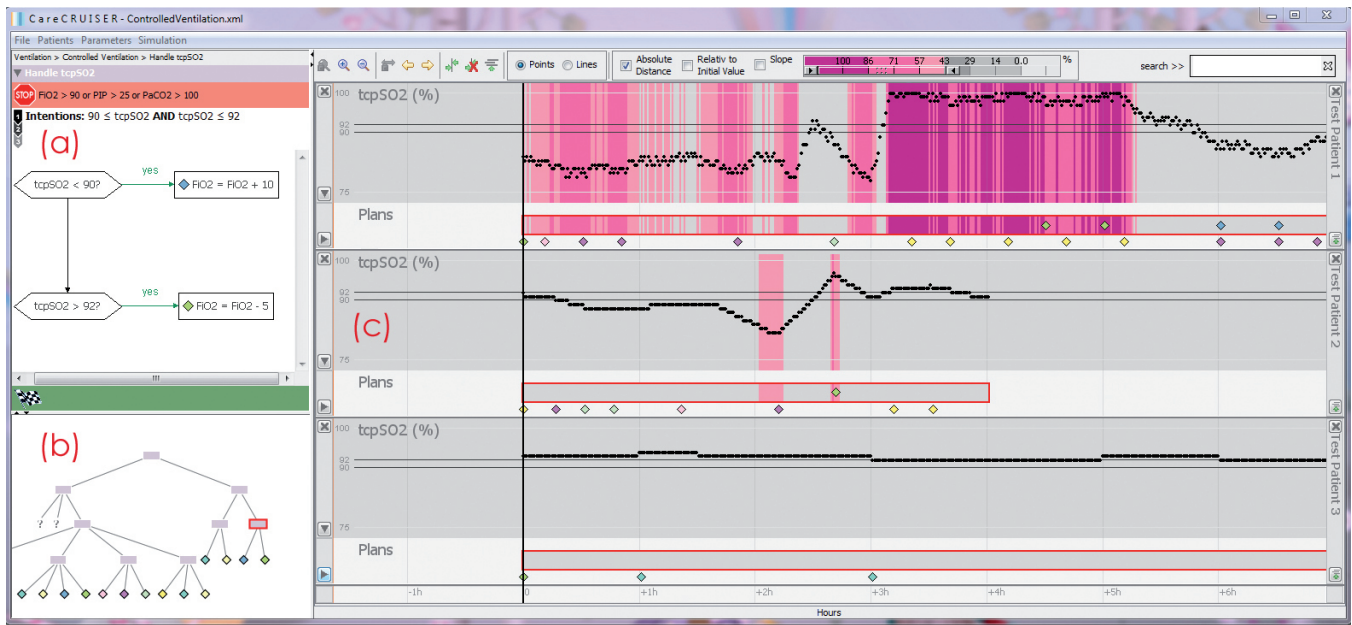


Figure 1: UI of the CareCruiser prototype. The logical view (a) communicates the logical structure of treatment plan execution by means of a flowchart-like representation [3]. The lower left part (b) displays a tree graph to visualize the hierarchical structure of treatment plans and sub-plans; the time-oriented 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. This screenshot shows one treatment plan that has been applied on three different patients (aligned vertically for comparison). The charts and treatment plans are colored according to the color scheme of the parameter values' distance to the intended value. Selecting ranges with big distance to the intended value with the range slider draws attention to critical cases and brings out the differences between the conditions of the three patients.

CareVis already provides three different views to communicate specific information, namely the logical view to visualize the logics of treatment plans, a view to show the hierarchical structure of these plans, and the temporal view to represent complex time constraints of treatment plans. The temporal view also allows for zooming and navigating along the time line (for detailed information we refer to [3]).

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

4.1 Clinical Actions

For assessing the effects of clinical actions the user needs to know when, how often, and in which context these actions have been applied. To visually represent these actions we chose diamonds for two reasons (see Figure 2):

1. Clinical actions can take place at one single point in time. Thus, the visual representation must allow for the identification of this exact point in time; this is indicated by the peak of the diamond.
2. Moreover, diamonds have a body which makes them more visible than simple vertical lines (which is especially true for actions at the start and end of a treatment plan).

In addition, we extended the representation with respect to clinical actions that take more time to be carried out. The temporal bounds of these actions are indicated by whiskers that span the duration of the action and allow for the identification of the exact time of its

start and end. Representing the start and end time by the width of the diamond would lead to line-thin diamonds for actions that take place at one single point in time, and thus, it would decrease the visibility of these actions. Similar to the representation of parallel plans, actions which are executed in parallel are aligned below each other. We assign a color (based on the qualitative color palettes proposed by Harrower and Brewer [12]) to each clinical action to facilitate their identification, for instance, two or more blue diamonds in a temporal treatment plan glyph indicate two or more instances of applying one clinical action (e.g., the application of a certain drug), in context of executing this treatment plan. Moreover, the color of clinical actions facilitates the coordination between the different views.




Figure 2: The gray rectangle represents a treatment plan along a time-axis containing diverse clinical actions. These clinical actions are represented by diamonds. The peak of the diamond indicates the exact point in time when the actions was 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 were applied to the patient but are not part of the treatment plan are laid out below the plan body.

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

sufficient for our test datasets.

4.2 Effects of Applied Treatment Plans

In order to determine success or failure of an applied treatment plan we take advantage of the intentions specified for this plan (intentions represent goals and intended effects of the treatment plan).

To visually encode how well an applied treatment plan succeeds in reaching the specified intentions and to make effects of applied clinical actions stand out more clearly, we use color-coding to highlight interesting events (i.e., distance information and slope) of the patient’s parameter development. We went for color-coded information highlighting (instead of, for instance, distance plots or plots of the first derivation) in order to maintain raw data charts without modifications or abstractions. This is important for physicians as they are used to assess a patient’s condition by these raw values. Besides the highlighting of color-coded information, the step-wise investigation of effects of applied treatment plans and clinical actions is supported by the following interactive means: filtering for interesting events of parameter values’ development, aligning plans or actions below each other to ease comparison (similar to the means for aligning clinical events in LifeLines2 [27]), and selecting regions of interest by means of a focus window. In the following subsections we give a detailed outline of these means.


4.2.1 Color-Coded Distance Information and Slope

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


Distance to Intended Value The distance of the patient’s parameter value (which is affected by the treatment plan) from the intended value is an important information to immediately identify critical values. We have enhanced the visibility of this information by using a color scheme which maps the relative distance of parameter values from the most deviant value within the range of possible parameter values (the minimum or maximum value of the parameter) to the intended value (e.g., 90-92% oxygen saturation) on a sequential color scale [12] from dark magenta to light magenta (see Figure 3(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 1 shows one treatment plan that has been applied on three different patients, which are vertically aligned. The charts and treatment plans are colored according to the patients’ parameter values’ distance to the intended value. We selected ranges with a big distance to the intended value with the range slider, which draws attention to critical cases. In doing so, the visualization emphasizes


the bad $tcpSO_2$ values of the first patient. In contrast, the minor changes in the $tcpSO_2$ curve of the third patient are close to the intended value range and do not draw off the user’s attention.



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



(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. For a more robust coloring we take the mean value of seven data points to compute the slope.

Figure 3: Comparing different modes of highlighting the effects of applied treatment plans. Each of them helps to reveal different aspects.

Progress from Initial Value The main criterion for the appraisal of treatment plans is how well the plan succeeds in achieving the intended effects. To visually encode this information, we consider the relative progress from the initial value (e.g., the oxygen saturation at the start of applying the treatment). The initial value is associated with 0% success and the intended value with 100% success of the treatment. In this way, success is expressed by the extent to which the oxygen saturation progresses from the initial value 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 effects.

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

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

To visually communicate the different qualities of progress of applied treatment plans we use a diverging color scale [12] from blue over white to red. We map the success values from 0% (value at plan start) to 100% (intended value) on colors from white to blue. In case the patient's values move into the opposite direction, the values from 0% to -100% are mapped on a color range from white to red. In this way, blue means good, red means bad, and, again, more saturated colors represent more extreme values (see Figure 3(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. So far, this is only an enhancement of information that is also given by the line chart. However, when it comes to getting an immediate impression of a patient's condition that depends on several different parameters, or to assessing the success of treatment plans which may intend to affect two or more parameters (e.g., 'stabilize the patient's $tcpSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 mmHg'), things get difficult: data can no longer be retrieved easily from the raw individual parameter curves. The overall success of the plan as well as the overall condition of the patient have to be determined. To this end, we stick to the rules of the infinite-valued Gödel logic [9]. According to these rules, the overall success of a plan that intends to achieve an effect on parameter A and an effect on parameter B is determined by the minimum of both success values (i.e., success (A and B) = minimum (success (A), success (B))). In turn, the overall success of a plan that intends to achieve a given effect on parameter A or an effect on parameter B is determined by the maximum of both values (i.e., success (A or B) = maximum (success (A), success (B))).

Figure 5 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 $tcpSO_2$ between 90% and 92% and the patient's PCO_2 between 40 and 60 mmHg', thus, the overall success of treatment plan execution at a given point in time is determined by the minimum of the success value for stabilizing the patient's $tcpSO_2$ 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 $tcpSO_2$ values are still too high. This results in a quite low overall success rate of the plan.

Slope Visually enhancing the slope of the parameter chart, i.e., if the parameter values rise or drop, reveals some interesting insights as well. In particular, it eases the identification of an immediate correlation of applied clinical actions and changes of the parameter value. In contrast to the other two highlighting modes, we don't encode 'good' and 'bad' changes (in terms of moving toward the intended value) in the parameter chart. This is why we decided for a diverging color scheme [12] that is as neutral as possible: we represent negative slopes with turquoise highlighting and positive slopes with brown highlighting (see Figure 3(c)).

4.2.2 Filtering Color-Coded Information

The user may be interested in only specific curve events, for instance, critical values or extreme drops of the curve. Thus, we provide a range slider to filter for such events of interest which makes them stand out immediately (see Figure 4).

4.2.3 Aligning Plans or Actions

To facilitate comparison between the effects one treatment plan had on two different patients or the effects of two alternative treatment plans we provide a tool to align selected treatment plans below each other together with their parameter curves (see Figure 1). The color




Figure 4: Selecting only extreme drops of the parameter chart with the range slider makes the coloring of minimal drops and rises of the chart disappear.

highlighting of these plans gives an immediate impression of differences in how well the plans have worked.

Similar to the comparison of treatment plans, we provide means to compare all instances of applying a specific clinical action. This allows the user to investigate how well this specific clinical action works (e.g., the application of a certain drug). The effects of applying this action once may not be representative; also, looking for all applications of this action and subsequently comparing them without visual assistance can be quite a laborious task. Hence, when selecting a clinical action of interest, all instances of applying this action are highlighted; together with the parameter chart they are aligned below each other to facilitate the comparison of the effects the applications of this action have had on the patient's parameters (see Figure 6).

4.2.4 Focus Window

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

Example: Discovering Less Efficient Treatment Choices
When assessing our new methods with our test datasets we could reveal an interesting pattern. The intention of the 'Handle $tcpSO_2$ ' plan is to stabilize the patient's $tcpSO_2$ between 90% and 92% by adjusting the supply of FiO_2 . We selected the clinical action 'reduce the FiO_2 when the patient's oxygen saturation rises above 92%' and vertically aligned all applications of this action. When highlighting the slope of the patient parameter $tcpSO_2$ and filtering for negative slopes (i.e., the oxygen saturation drops), we can see that almost every time this action was applied the oxygen saturation dropped after some delay. In the meantime the action was applied again which caused another drop of the oxygen saturation below 90%. Therefore, the complementary action had to be applied ('increase the FiO_2 when the patient's oxygen saturation drops below 90%'). Helping physicians to reveal this behavior leads to more effective treatment, a reduction of treatment costs, and a better stabilization of the patient (see Figure 6).

4.3 Investigation of Multiple Patients

The CareCruiser prototype allows for the investigation of two or more patients at the same time. The available display height is equally distributed to the patient facets (see Figure 1). However, when the patient facets reach a defined minimum height, the view gets scroll-able. To this end, collapsed patient facets allow for investigating a bigger number of patients at the same time (see Fig-

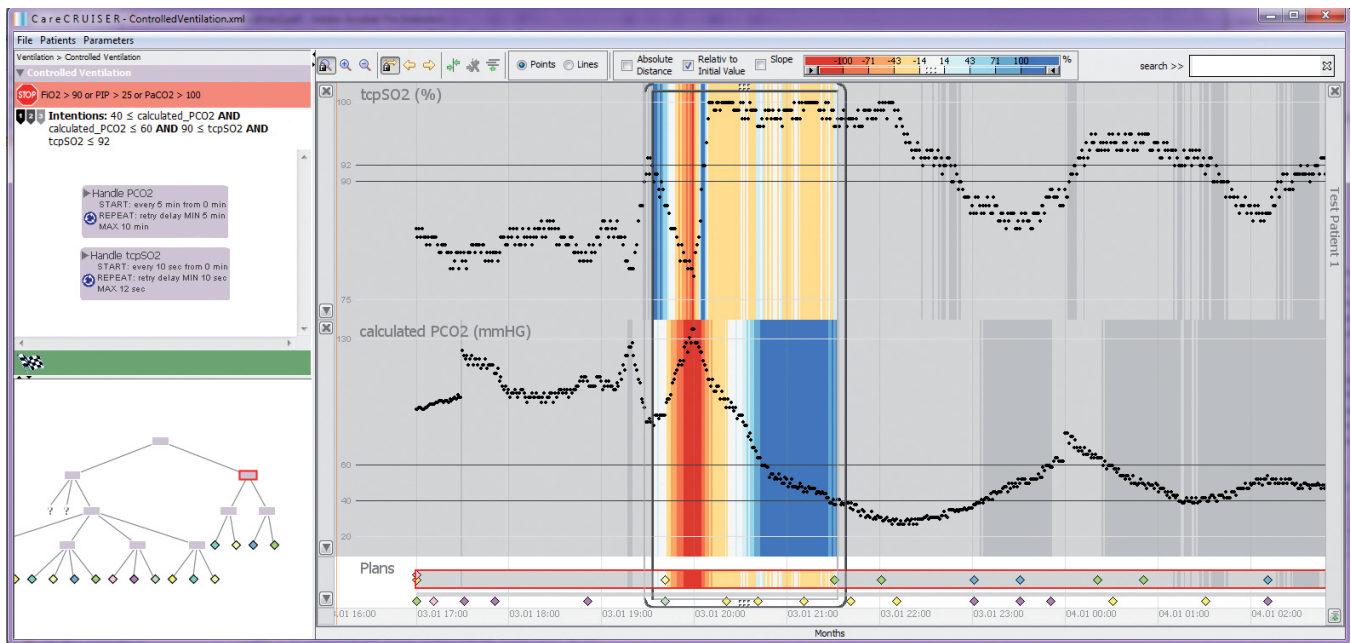


Figure 5: Highlighting the progress of parameter values from the initial value toward the intended value, results in getting an immediate idea of success or failure of the applied treatments. The overall success of executing the selected treatment plan at a given point in time is determined by the minimum of the success value for stabilizing the patient's $tcpSO_2$ and the success value for stabilizing the patient's PCO_2 . The low success values for stabilizing the patient's $tcpSO_2$ lead to a low overall success for the selected treatment plan.

ure 7). These facets can be expanded, collapsed or closed individually or collectively. When moving the mouse over the right border of the patient facet a pop-up window provides demographic information about the patient such as name, age, gender, etc.

5 EVALUATION AND FURTHER WORK

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

To ensure the usability of the prototype we decided for a heuristic usability evaluation according to Forsell and Johansson [7]. They presented a best practice set of 10 heuristics out of 63 heuristics (from 6 earlier published heuristic sets). This new set is especially tailored to the evaluation of common and important usability problems in Information Visualization techniques.

It is commonly assumed that three to five expert evaluators are sufficient for a heuristic usability evaluation [14]. Thus, we conducted the study with two female and two male evaluators which have a degree in computer science. In the context of their research, all of them have gained practical and theoretical experience in the field of usability engineering.

To ensure unbiased evaluations, we performed separate testing sessions with each evaluator. An observer was present at each testing session to answer questions about the domain and to give hints when the evaluator was in trouble (which was not necessary at any time of the four testing sessions). Each evaluator went through the visualization two times. The first round was aimed at getting a feeling for the flow and the general scope of the visualization. In the second round the evaluator was supposed to focus on visual and interactive interface elements with respect to the given list of heuristics [7]. Each evaluator was asked to solve a list of tasks by means of the CareCruiser prototype (1. to find out how the patient's con-

dition changes during the execution of the treatment plan; 2. to find out which clinical actions were applied to the patient in the context of treatment plan execution and when these actions were applied; 3. to find out which effects the clinical actions have on the patient's condition; 4. to identify critical parameter values in the course of treatment plan execution) and to note down the found problems with reference to the violated usability principles. Moreover, the evaluators were asked to rate the severity of the problem (1 representing the lowest severity and 5 the highest).

This evaluation resulted in a list of 32 usability problems. The majority of these problems relate to the two usability principles 'orientation and help' and 'consistency'. 'Orientation and help' deals with means such as support to control levels of details, redo/undo of actions, and representing additional information, while 'consistency' refers to the way design choices are maintained in similar contexts, and are different when applied to different contexts [7].

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

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

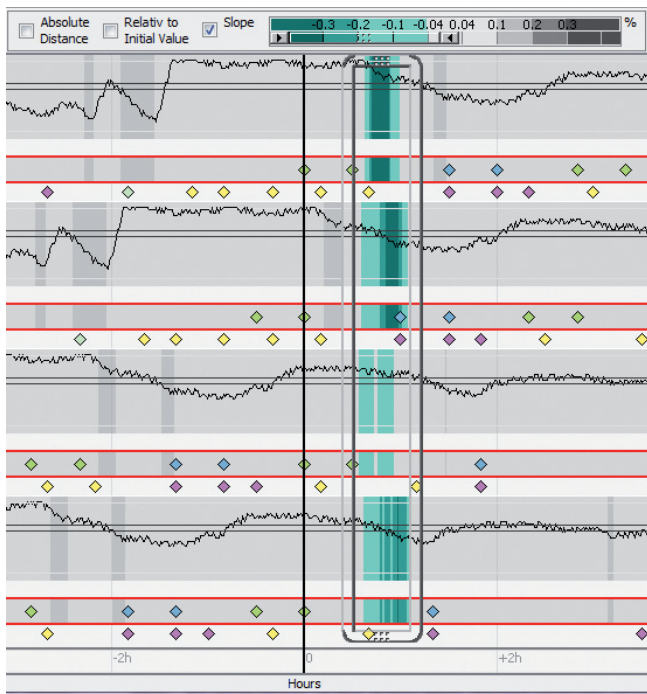


Figure 6: Revealing the delayed drop of the patient's $tcpSO_2$ when applying the clinical action 'reduce the FiO_2 when $tcpSO_2$ rises above 92%'. All applications of this action were aligned vertically along the black line; the negative slopes of $tcpSO_2$ were highlighted. Dragging the focus window over the time span after applying the action reveals a vertical turquoise pattern (drops of the $tcpSO_2$ curve) with some delay to the application of the action.

these treatment plans.

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



Figure 7: Collapsed patient facets. We have aligned a given treatment plan that has been applied to four different patients. When highlighting the parameter values' progress from the initial value and filtering for parameter values within the intended range (blue) one can still tell that the treatment plan succeeded in stabilizing the values of the first and third patient while failed for the second and fourth patient.

adequate.

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

However, we will carefully consider all found problems in order to improve the CareCruiser prototype and we plan to conduct additional scenario-based evaluations with physicians.

In the near future we will also investigate other possible areas of applying our visual encodings. In general, our visual encodings could help to get new insights in any domain where causes and effects are to be investigated in a temporal context. This may hold true for economic context (such as certain economic incidents and their effects on currency rates), for public relations context (such as the effects of press reports and public relations on the findings of voter surveys), academic context (such as the effects of introducing new structural or curricular elements in education on quantity and quality of graduates and/or scientific output in different countries), a.s.o. We plan to investigate the various possible areas of applying these visual encodings in the near future.

6 CONCLUSION

CareCruiser is based on the architecture of CareVis [3], but is a conceptual extension thereof. First, CareCruiser takes the process information (i.e., the specified intentions of treatment plans) into account to visualize the processes' effects instead of pure data visualizations. Second, CareCruiser provides various interaction methods to explore the data, processes, and effects interactively. Third, CareCruiser offers different views about these effects. Fourth, CareCruiser allows for evaluating multiple patients in parallel. It has been developed to assist physicians and medical researchers both with improving treatment plans and with the optimization in applying them. By now, some limitations of the visualization are that it supports only continually measured, quantitative data, it does not support more detailed semantic information about each patient parameter (as described in Section 5), and it does not support treatment plan goals like maintaining a certain condition after the treatment has finished. Being aware of these limitations, we believe that CareCruiser can easily be extended to provide also this functionality.

While there are a great number of sophisticated tools available that focus on individual aspects of the application of computer-executable treatment plans, our prototype takes the available methods a decisive step further by providing a working environment that offers an intuitive comprehension and grasp of aspects being of interest in the context of evaluating and executing treatment plans. CareCruiser offers three main advantages over other tools in the field:

1. Despite the wealth of sophisticated information yielded by CareCruiser, it is easy to work with, due to its self-explanatory way of handling and presenting data. We have conducted a usability evaluation and evaluated the visualization scenario-based with our cooperating physician. Thus, it may readily be used both in research and in monitoring environments (e.g., hospitals).
2. CareCruiser's comprehensive integration of functions for a step-wise interactive exploration of effects of applied treat-

ment plans on patient parameters to support the revelation of new insights and the generation of hypotheses. These means are designed with special respect to support the assessment of success or failure of plans, and thus, to support the identification of the best working therapies. There are:

- Means to highlight and filter specific events of parameter developments,
 - Means to align treatment plans or clinical actions together with the corresponding patient parameter values to ease comparison,
 - Means to move the focus along the time line to support the detection of interesting event patterns while maintaining context information.
3. Of particular interest may be the provision to visually compare both details and the results of treatment applied to a great number of patients. Due to its presentation and filtering options, CareCruiser can be fine-tuned to offer much contextual evidence at a glance, thereby being of substantial help in the formulation and/or identification of the best working therapies.

ACKNOWLEDGEMENTS

We wish to thank Christian Popow, Dept. of Child and Adolescent Psychiatry, Medical University of Vienna, for his cooperation and his helpful suggestions and the four evaluators for their time and thoroughly investigation of the CareCruiser prototype. Moreover, the research leading to these results 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).

REFERENCES

- [1] W. Aigner. Guideline overview tool (GOT), Asgaard-TR-2001-4. Technical report, Vienna University of Technology, Institute of Software Technology and Interactive Systems, Vienna, Austria, 2001.
- [2] W. Aigner, K. Kaiser, and S. Miksch. *Visualization Techniques to Support Authoring, Execution, and Maintenance of Clinical Guidelines*, volume 139: Studies in Health Technology and Informatics, pages 140–159. 2008.
- [3] W. Aigner and S. Miksch. CareVis: Integrated visualization of computerized protocols and temporal patient data. *Artif Intell Med*, 37(3):203–218, May 2006.
- [4] R. Bade, S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. In *CHI '04: Proc. of the SIGCHI conference on human factors in computing systems*, pages 105–112, 2004.
- [5] 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.
- [6] M. J. Field and K. N. Lohr, editors. *Clinical Practice Guidelines: Directions for a New Program*. National Academies Press, Institute of Medicine, Washington DC, 1990. <http://www.nap.edu/books/0309043468/html/> (last accessed: Sept. 21, 2010).
- [7] C. Forsell and J. Johansson. An heuristic set for evaluation in information visualization. In *AVI '10: Proceedings of the International Conference on Advanced Visual Interfaces*, pages 199–206, New York, NY, USA, 2010. ACM.
- [8] 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, 2002.
- [9] K. Gödel. Zum intuitionistischen Aussagenkalkül (On the intuitionistic propositional calculus). *Anzeiger der Akademie der Wissenschaften in Wien*, 69:65–66, 1932.
- [10] T. Gschwandtner. Visualization of patient data and treatment plans: A survey. Asgaard-TR-2009-1. Technical report, Vienna University of Technology, Institute of Software Technology and Interactive Systems, 2009.
- [11] D. C. Hadorn. Use of algorithms in clinical practice guideline development: Methodology perspectives. *Clinical Practice Guideline Development: Methodology Perspectives*, 9(95):93–104, Jan. 1995.
- [12] M. Harrower and C. Brewer. Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, December 2003.
- [13] 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.
- [14] A. Holzinger. Usability engineering methods for software developers. *Communications of the ACM*, 48(1):71–74, 2005.
- [15] W. Horn, C. Popow, and L. Unterasser. Metaphor graphics to visualize ICU data over time. In *Proc. of Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-98), Workshop Notes of the ECAI-98 Workshop*, pages 76–81, Aug. 1998.
- [16] 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'06)*, pages 404–411. ACM, May 2006.
- [17] D. Klimov and Y. Shahar. A framework for intelligent visualization of multiple time-oriented medical records. In *Proc. of the AMIA Annual Symposium (AMIA 2005)*, pages 405–409, 2005.
- [18] R. Kosara and S. Miksch. Metaphors of movement: A visualization and user interface for time-oriented, skeletal plans. *Artif Intell Med, Special Issue: Information Visualization in Medicine*, 22(2):111–131, May 2001.
- [19] P. McLachlan, T. Munzner, E. Koutsofios, and S. North. LiveRAC - interactive visual exploration of system management time-series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'08)*, pages 1483–1492. ACM, April 2008.
- [20] C. Plaisant, S. J. Lam, B. Shneiderman, M. S. Smith, D. H. 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. In *Proc. of the AMIA Annual Symposium (AMIA 2008)*, pages 601–605, 2008.
- [21] 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 *Proc. of the AMIA Annual Symposium (AMIA 1998)*, pages 76–80, 1998.
- [22] S. M. Powsner and E. R. Tufte. Graphical summary of patient status. *The Lancet*, 344:386–389, 1994.
- [23] S. Quaglinia, M. Stefanelli, G. Lanzola, V. Caporusso, and S. Panzarasa. Flexible guideline-based patient careflow systems. *Artif Intell Med*, 22(1):65–80, April 2001.
- [24] 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. *Artif Intell Med*, 38(2):115–135, Oct. 2006.
- [25] 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.
- [26] P. Terenziani, S. Montani, A. Bottrighi, M. Torchio, G. Molino, and G. Correndo. The GLARE approach to clinical guidelines: Main features. *Stud Health Technol Inform*, 101:162–166, 2004.
- [27] 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 *Proc. of the ACM Conference on Human Factors in Computing Systems (CHI 2008)*, pages 457–466, Jan. 2008.