

Design and Evaluation of an Interactive Visualization of Therapy Plans and Patient Data

Theresia Gschwandtner, Wolfgang Aigner, Katharina Kaiser, Silvia Miksch, and Andreas Seyfang
Institute of Software Technology & Interactive Systems (ISIS)
Vienna University of Technology
Favoritenstrasse 9-11/188
A-1040 Vienna
Austria
ieg.ifs.tuwien.ac.at
{gschwandtner, aigner, kaiser, miksch, seyfang}@ifs.tuwien.ac.at

CareCruiser is a system designed to visualize the effects of applying clinical treatment plans and to support the exploration of the effects that the treatments have on the patient's condition. While considerable progress has been made in recent years concerning the visualization methods to support computer-executable clinical treatment plans, our prototype succeeds in combining and enhancing the possibilities offered by the features of other systems. To ensure the quality of our visualization (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 the exploration features of CareCruiser.

Case study, usability evaluation, medical information visualization, clinical practice guidelines

1. INTRODUCTION

Clinical practice guidelines (CPGs) are text documents that provide treatment recommendations for specific clinical situations according to best medical evidence. Their ability to improve clinical care is well recognized (see Field and Lohr (1990)). While physicians are expected to adhere to these guidelines, their actual use strongly depends on the immediate availability of these recommendations (see Elkin et al. (2000)). Thus, making the CPG's logic easily accessible (i.e., making the guidelines computer-executable and providing effective visualizations to communicate the recommendations) increases their acceptability.

In order to make the CPG's recommendations easily accessible we have developed an interactive information visualization system (introduced in Gschwandtner et al. (2011)). This visualization not only effectively communicates the CPG's logic, but also facilitates the exploration of the effects of applied treatment plans, which helps to assess clinical situations. The investigation of the actual

effects of previously applied treatment plans and clinical actions on the patient's condition helps to optimize the choice of future treatment actions (CPGs sometimes contain tolerance in when and how often given recommendations are to be applied). Our visualization, called CareCruiser, enhances existing visualization techniques (see Aigner and Miksch (2006)). In Gschwandtner et al. (2011) we presented the visual encodings and the interactive features of CareCruiser. However, a careful evaluation of the design and the features of the prototype is of major importance to ensure its quality. Hence, in this paper, we give a detailed outline of the evaluation process and its outcome.

First of all, we give an outline of related work in Section 2. We continue with a short description of the visual encodings and the interactive features of our visualization in Section 3. In Section 4 we outline the evaluation of this visualization – a heuristic usability evaluation and user testing. Finally, we sum up the main results of our work in Section 5.

2. RELATED WORK

There is still a noticeable shortage of projects dealing with the visualization of applied treatment plans in combination with patient data. Merely visualizing either treatment plans or patient data is encountered more often, as for instance in the following projects (for a detailed description we refer to Aigner et al. (2008)). There are a number of visualizations such as Protégé (see Gennari et al. (2002)) and GUIDE (see Quaglinia et al. (2001)) that provide flow-chart-like representations of clinical algorithms. They are suited to show the execution sequence of plans, but they leave out the aspects of temporal constraints of this execution sequence of plans. A more sophisticated approach is executed by AsbruView (see Kosara and Miksch (2001)), as AsbruView can visualize many relevant plan characteristics; however, it does not visualize patient data.

Patient data, on the other hand, is the main theme of visualizations such as Graphical Summary of Patient Status (see Powsner and Tufte (1994)) which uses small repeated graphs to draw a comprehensive picture of a patient's condition. LifeLines (see Plaisant et al. (1998)) represents specific aspects of a patient record by horizontal lines on a time scale, LifeLines2 (see Wang et al. (2008)) visualizes subsets of medical records from multiple patients, and PatternFinder (see Plaisant et al. (2008)) is used to query, display, and align special events of patient records to reveal interesting patterns. KNAVE II (see Shahar et al. (2006)) provides multiple features, like different abstraction levels, absolute and relative time scales, different granularities of time, and additional statistics to support the exploration of patient parameters. VIE-VISU (see Horn et al. (1998)) uses glyphs to represent very specific characteristics of neonatal intensive care data. Gravi++ (see Hinum et al. (2005)) uses animation and traces to visualize the change of the patient's condition over time. They all deal with various aspects of depicting a patient's condition, a dominant aim being the revelation of treatment relevant patterns. However, none of these approaches is aimed at communicating characteristics of a given treatment plan application.

The requirements of combining the visualization of any applied treatment plans with corresponding patient data is met by very few systems: the Guideline Overview Tool (see Aigner (2001)) combines parameter charts with an extension of LifeLines to display basic treatment plan characteristics, Midgaard (see Bade et al. (2004)) represents detailed patient data and complex plan characteristics, and CareVis (see Aigner and Miksch (2006)) visualizes different aspects of treatment

plans by means of multiple views in combination with patient parameter charts.

Outside the field of patient data and medical treatment plans, there are two programs to be mentioned that use visualization techniques related to the visual encodings of CareCruiser: LiveRAC (see McLachlan et al. (2008)) arranges huge numbers of data in a matrix and offers features such as the interactive arranging of rows and columns and semantic zooming, all of which allow for investigating these data at multiple levels of detail. The Line Graph Explorer (see Kincaid and Lam (2006)) concentrates on visualizing large amounts of data in a limited display space by using not only different colors but also color saturation and luminance. Although our own approach to depict value ranges of multiple line charts is similar to that of the Line Graph Explorer, our color codings go one step further by including such semantic information as treatment plan goals and parameter progress. We use color to highlight and filter for specific events of interest and to visually relate clinical actions to line chart events.

Of the above mentioned programs dealing with the visualization of either treatment plans or patient data, none provides insight into the effects of specific events on a patient's condition – for instance, the effects of the repeated administration of a certain drug; neither can these effects be examined interactively with their help.

3. DESIGN

In the following subsections we briefly describe the visual encodings and interactive features of CareCruiser. For a detailed description of CareCruiser we refer to Gschwandtner et al. (2011).

CareVis (see Aigner and Miksch (2006)), the system on which CareCruiser is based, offers three views to show specific information (see Figure 1): (a) the logical view visualizing the logics of treatment plans, (b) the hierarchical view showing the hierarchical structure of these plans, and (c) the temporal view depicting time constraints of treatment plans. We have extended the temporal view to provide several features to support a step-wise interactive exploration of the patient's condition and effects of applied treatment plans:

- Vertically aligning treatment plans to ease comparison of the effects of different plans or the effects on different patients.
- Vertically aligning all instances of a given clinical action to investigate the effects of this action (see Figure 3).

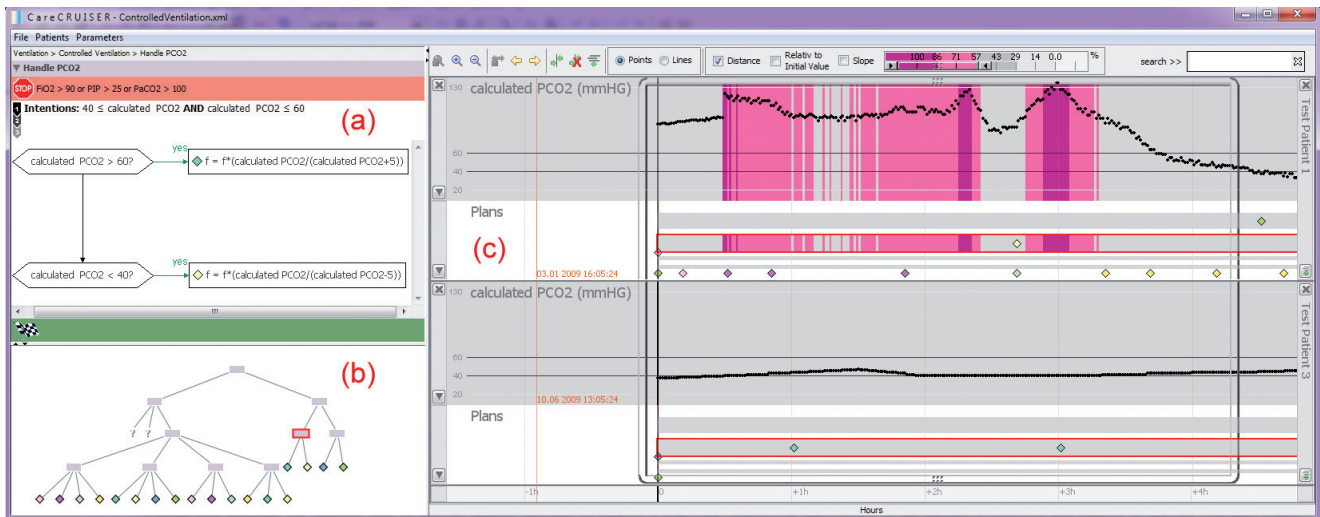


Figure 1: UI of the CareCruiser prototype. The logical view (a) communicates the logical structure of treatment plan execution by means of a flowchart-like representation (see Aigner and Miksch (2006)). The lower left part (b) displays a tree graph to visualize the hierarchical structure of treatment plans and sub-plans; the temporal view (c) focuses on the temporal-qualities of applied treatment plans, clinical actions, and patient parameters. One treatment plan that has been applied on two 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. Filtering for ranges with a large distance to the intended value (critical cases) using the range slider shows the differences between the conditions of the two patients.

- Color-coded highlighting of interesting events of the development of a patient parameters' values (see Figure 2).
- Filtering these highlighted events (see Figure 3).
- Providing *focus + context* techniques to support the detection of patterns (see Figure 3).
- Comparing two or more patients simultaneously (see Figure 1).

Our extensions thus serve a two-fold purpose: with their help, it is easy to grasp the effects of applied treatment plans on patients by using color-coded distance information and color-coded sloping of parameter curves, and also by being able to compare multiple patients simultaneously. The visualization and interactive features described in the following can be used to explore the clinical situation of one patient in detail or to compare the individuals of a group of patients. The former allows a detailed analysis of individual reactions to clinical actions of one single patient and consequently the optimization of a specific treatment plan. The latter provides on-the-spot comparison of either treatment variations or treatment effects on different patients.

3.1. Visual Encoding

Our color coding system brings to light a number of treatment-relevant aspects that otherwise had not been easy to grasp (see Figure 2). They are oriented


towards the patient parameters in combination with the applied treatment plans and show

- (a) the distance of the parameter value from the intended value,
- (b) the progress of the parameter value from the initial value, and
- (c) the rise and fall of the parameter values.


These color schemes (based on the diverging color palettes proposed by Harrower and Brewer (2003)) encode sections of the parameter chart referring to specific events, for instance, the downward tilt of a parameter curve, or parameter values far off the intended value range (i.e., critical parameter values). It is thus very easy to compare different treatment plans, as it is color that allows us to intuitively assess data details otherwise revealed only by more or less meticulous scrutiny of either mathematical or graphical information.

3.2. Interactive Features


CareCruiser's interactive features comprise zooming and navigating along time, a search function to filter and highlight plans and clinical actions, and details on demand for treatment plans, clinical actions, and patients' demographics. Moreover, they contain a range slider to filter the color highlighting of curve events that a user might be interested in. Then there is a tool either to align selected treatment plans below each other, together with their parameter charts, or to align all instances of applying a given



(a) The **distance to the intended value** color scheme helps physicians to identify critical values at the first sight. The range of intended values is indicated by the two dark horizontal lines (dark magenta: extreme values, light magenta: inside the intended value range).



(b) The **progress from the initial value** (relative to the initial value when the treatment plan was started) color scheme shows to what extent the applied treatment plan has the intended effect on the patient's condition (white: start value, dark blue: intended value, dark red: departure from the intended value).



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

Figure 2: Different modes of color-coding the effects of applied treatment plans. Each of them helps to reveal different aspects.

clinical action (see Figure 3). The latter makes it easy to find out how well a specific clinical action works (e.g., the application of a certain drug). We also provide a focus window to gray out the color information outside its borders; thus the window can be focused on any time section of particular interest (see Figure 3). The window width (i.e., the length of the time span in focus) can be varied. Of course the position of the window along the time line can be varied as well.

4. EVALUATION

We took several actions 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 visualization 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 methods. Finally, (3) we tested the system with two target users (i.e., medical experts) to gain insight 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).

4.1. Usability Evaluation

In order to improve the usability of CareCruiser, we conducted a heuristic usability evaluation according to Forsell and Johansson (2010). The reason why we decided to use this set of heuristics instead of the classical heuristics according to Nielsen (1994) is that this new set by Forsell and Johansson is especially tailored to the evaluation of important usability problems in Information Visualization techniques. Forsell and Johansson present a best practice set of 10 heuristics¹ out of 63 heuristics (from 6 earlier published heuristic sets).

We conducted the study with four evaluators according to the common assumption that three to five evaluators are sufficient for a heuristic evaluation (see Holzinger (2005)). 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 Forsell and Johansson (2010)). 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).

¹a complete list of these heuristics can be found on the project page: <http://ieg.ifs.tuwien.ac.at/projects/carecruiser/>

32 usability problems were encountered in the course of this usability evaluation. Most frequent among them were the two usability principles '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 Forsell and Johansson (2010)).

Most of the problems found were not severe. There have been found, however, a couple of problems that require a little bit more work 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. We will correct all of the found problems in the near future.

4.2. User Testing

We presented the CareCruiser prototype to two physicians, as physicians are the 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 was formerly 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 visualization and its features, the physicians used the visualization 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, Miles and Huberman (1994); Yin (1994); Neuman (2000)). 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 are captured, the interview was audio-recorded. Subsequently, the interview was transcribed into text and the transcripts were reviewed by the interview subjects 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 and 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., 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's 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 (translated):

'The coloring [of the parameter's distance to the intended value] seems to improve the visibility of the patient's condition.'²

²translated from German by the authors: 'Die Farben scheinen die Übersicht über den Patientenzustand zu verbessern.'

‘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.’³

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’s 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 window to inspect a short span of the time following the application of the action. The physician reasoned that the patient did not receive the 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 it 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 (see Figure 3). The physician also agreed that the visualization 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 (translated):

‘It is a very feasible solution with interesting features to analyze this kind of data.’⁴

‘The strengths of the visualization are, that it is clear, easy to work with [...] - intuitive, beautiful.’⁵

³translated from German by the authors: ‘Bei dieser Visualisierung ist sehr schön nachvollziehbar welche Aktionen warum gesetzt wurden.’

⁴translated from German by the authors: ‘Das ist eine sehr praktikable Lösung diese Daten zu analysieren mit interessanten Features.’

⁵translated from German by the authors: ‘Die Stärken sehe ich darin, dass die Visualisierung übersichtlich ist, dass sie handlich ist [...] intuitiv, wunderschön.’

‘I believe [by means of this tool] you can very well demonstrate the clinical situation to students.’⁶

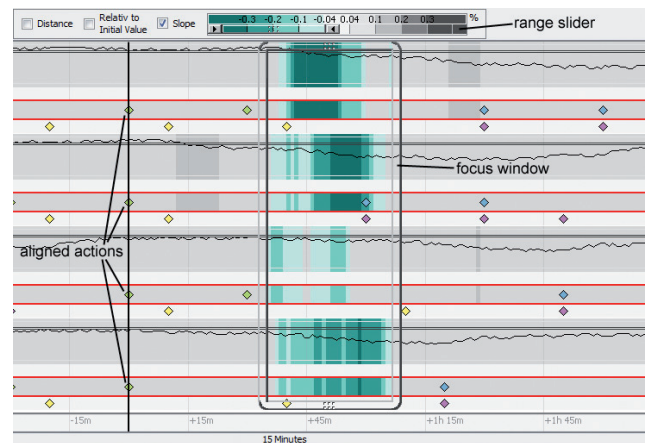


Figure 3: All applications of a clinical action that is supposed to drop the displayed parameter were aligned vertically along the black line; the negative slopes of the parameter curve were highlighted in turquoise. Dragging the focus window over the time span after applying the action reveals a vertical turquoise pattern (drops of the parameter curve) with some delay to the application of the action.

4.3. Outcome

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. It took both of them very little time to learn how to use the system. In approximately 15 minutes they were able to navigate between views, add and remove patient parameters, apply different color-highlightings, and align clinical actions. While interacting with CareCruiser, they stumbled 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 according to these findings. Both physicians eagerly used CareCruiser’s interactive features to explore the given clinical situation. When being asked about shortcomings of the visualization 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

⁶translated 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 [...]’

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).

Subsequently to the interaction phase, we conducted semi-structured interviews with 13 questions⁷, 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 effects on a patient's condition?
3. Does CareCruiser help to detect unexpected effects of clinical actions and thus help to optimize treatment choices?

Both physicians gave affirmative answers to practically all of these questions. 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 established kind of acoustic alarm would do a better job. However, replacing acoustic alarms in critical situations has never been the intention of our visualization. Both physicians gave positive answers to all the questions relating to the three main categories mentioned above. They said that CareCruiser facilitates capturing the course (i.e., positive and negative effects) of an applied treatment; they affirmed that CareCruiser provides all information (about the patient and the treatment plan) necessary to assess the clinical situation, they stated that CareCruiser helps to detect unexpected patient reactions on applied clinical actions; and they were sure that the insights gained with CareCruiser help to optimize treatment. One physician stated that the visualization would also be of great value when a patient is passed to another physician (i.e., a shift handover), since it provides useful information about which clinical actions had positive effects on the patient's condition and which did not. General findings of the study include that the complex interaction techniques were 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 of derived values in combination with raw data values help to highlight specific events that could easily be overlooked otherwise.

⁷a complete list of these questions can be found on the project page: <http://ieg.ifs.tuwien.ac.at/projects/carecruiser/>

5. CONCLUSION

CareCruiser is a conceptual enhancement of the architecture of CareVis. We have designed CareCruiser to visualize treatment steps and their effects rather than pure data. The main advantage for clinicians lies in the 'added value' of our system. Firstly, it serves as an immediately comprehensible visual protocol of what was applied to a patient when and in what quantity, and how the patient reacted to it. Secondly, it provides features that help to explore in depth this cause and effect relationship interactively, and in doing so, to more readily gain insight into a complicated matter. This, then, points to the generalizing potential of our system: the decisive facts are the quantity of information that can be communicated by visualization, the clarity of this visualized information, and the ease of exploring and dealing with this information and drawing the correct conclusions from it.

To assure the quality of CareCruiser we (1) collaborated with a medical expert in the design phase. Moreover, we have subjected CareCruiser to an evaluation process concerning both (2) its usability and (3) its actual usefulness in a clinical environment. The outcome of the evaluation was very positive; apart from some minor deficiencies that can and will easily be corrected in the near future, the overall outcome of the evaluation showed that CareCruiser is a valuable enhancement of the instruments that have been available to physicians so far.

6. ACKNOWLEDGEMENTS

We wish to thank the two physicians, for their cooperation and careful evaluation of our system. Furthermore, we wish to thank the four usability evaluators for their time and thorough investigation of the CareCruiser prototype. 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).

7. REFERENCES

- Aigner, W. (2001) Guideline Overview Tool (GOT), Asgaard-TR-2001-4, Technical Report, Vienna University of Technology, Institute of Software Technology and Interactive Systems, Vienna, Austria.
- Aigner, W., Kaiser, K., and Miksch, S. (2008) Visualization Methods to Support Guideline-Based

- Care Management, in: *Computer-Based Medical Guidelines and Protocols: A Primer and Current Trends*, IOS Press, pp. 140-159.
- Aigner, W. and Miksch, S. (2006) CareVis: Integrated Visualization of Computerized Protocols and Temporal Patient Data, in: *Artif Intell Med*, 37(3):203-218.
- Bade, R., Schlechtweg, S., and Miksch, S. (2004) 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*, pp. 105-112.
- Elkin, P. L., Peleg, M., Lacson, R., Bernstam, E., Tu, S. W., Boxwala, A., Greenes, R., Shortliffe, E. H. (2000) Toward Standardization of Electronic Guideline Representation, in: *MD Computing*, 17(6):39-44.
- Field, M. J., Lohr, K. N. (ed.) (1990) *Clinical Practice Guidelines: Directions for a New Program*, National Academies Press, Institute of Medicine.
- Forsell, C. and Johansson, J. (2010) An Heuristic Set for Evaluation in Information Visualization, in: *Proc. of the International Conference on Advanced Visual Interfaces (AVI '10)*. ACM, pp. 199-206.
- Gennari, J. H., Musen, M. A., Ferguson, R. W., Grosso, W. E., Crubézy, M., Eriksson, H., Noy, N. F., and Tu, S. W. (2002) The Evolution of Protégé: An Environment for Knowledge-Based Systems Development, in: *International Journal of Human-Computer Studies*, 58:89-123.
- Gschwandtner, T., Aigner, W., Kaiser, K., Miksch, S., Seyfang, A. (2011) CareCruiser: Exploring and Visualizing Plans, Events, and Effects Interactively, in: *Proc. of the IEEE Pacific Visualization Symposium (PacificVis 2011)*, accepted.
- Harrower, M. and Brewer, C. (2003) ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps, in: *The Cartographic Journal*, 40(1):27-37.
- Hinum, K., Miksch, S., Aigner, W., Ohmann, S., Popow, C., Pohl, M., and Rester, M. (2005) Gravi++: Interactive Information Visualization of Highly Structured Temporal Data, in: *Journal of Universal Computer Science, Special Issue on Visual Data Mining*, 11:1792-1805.
- Holzinger, A. (2005) Usability Engineering Methods for Software Developers, in: *Communications of the ACM*, 48(1):71-74.
- Horn, W., Popow, C., and Unterasinger, L. (1998) 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*, pp. 76-81.
- Kincaid, R. and Lam, H. (2006) Line Graph Explorer: Scalable Display of Line Graphs Using Focus+Context, in: *Proc. of the Working Conference on Advanced Visual Interfaces (AVI'06)*, ACM, pp. 404-411.
- Kosara, R. and Miksch, S. (2001) Metaphors of Movement: A Visualization and User Interface for Time-Oriented, Skeletal Plans, in: *Artif Intell Med, Special Issue: Information Visualization in Medicine*, 22(2):111-131.
- McLachlan, P., Munzner, T., Koutsofios, E., and North, S. (2008) LiveRAC - Interactive Visual Exploration of System Management Time-Series Data, in: *Proc. of the SIGCHI Conference on Human Factors in Computing Systems (CHI'08)*, ACM, pp. 1483-1492.
- Miles, M. B. and Huberman, M. (1994) *Qualitative Data Analysis: An Expanded Sourcebook* (2nd edn). Sage Publications.
- Neuman, W. L. (2000) *Social Research Methods: Qualitative and Quantitative Approaches* (4th edn). Allyn And Bacon.
- Nielsen, J. (1994) Heuristic Evaluation, in: *Usability Inspection Methods*. John Wiley & Sons, pp. 25-62.
- Plaisant, C., Lam, S., J., Shneiderman, B., Smith, M., S., Roseman, D., H., Marchand, G., Gillam, M., Feied, C., Handler, J., and Rappaport, H. (2008) 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)*, pp. 601-605.
- Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., and Shneiderman, B. (1998) LifeLines: Using visualization to enhance navigation and analysis of patient records, in: *Proc. of the AMIA Annual Symposium (AMIA 1998)*, Hanley and Belfus, Inc., pp. 76-80.
- Powsner, S. M. and Tufte, E. R. (1994) Graphical Summary of Patient Status, in: *The Lancet*, 344:386-389.
- Quaglinia, S., Stefanelli, M., Lanzola, G., Caporusso, V., and Panzarasa, S. (2001) Flexible guideline-based patient careflow systems, in: *Artif Intell Med*, 22(1):65-80.
- Shahar, Y., Goren-Bar, D., Galperin, M., Boaz, D., and Tahan, G. (2006) Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions, in: *Artif Intell Med*, 38(2):115-135.
- Wang, T. D., Plaisant, C., Quinn, A., Stanchak, R., Shneiderman, B., and Murphy, S. (2008) 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)*, ACM, pp. 457-466.
- Yin, R. K. (1994) *Case study research. Design and methods* (2nd edn). Sage Publications.