

# Visual Analysis of Compliance with Clinical Guidelines

Peter Bodesinsky

Paolo Federico

Silvia Miksch

Institute of Software Technology and Interactive Systems  
Vienna University of Technology, Vienna, Austria  
peter.bodesinsky@gmx.at, {federico, miksch}@ifs.tuwien.ac.at

## ABSTRACT

Clinical guidelines provide recommendations in the form of applicable actions in a specific clinical context. Computer Interpretable Guidelines (CIG) aim to achieve guideline integration into clinical practice to increase health care quality. Analyzing the compliance with a CIG can facilitate the implementation and assist in the design of CIGs, but to help medical experts in the detection of patterns in the wealth of the data is a challenging task. We suggest an approach based on visual analytics, intertwining interactive visualization and automated data analysis i.e. analysis of compliance with a CIG. Our solution covers highlighting and abstraction for time-oriented patient parameters, and aggregation of repeatedly missing actions into intervals; in addition valid, invalid, and missing actions are represented visually. Furthermore, we discuss a case study showing how the applied techniques can assist in the detection of interesting patterns.

## Categories and Subject Descriptors

H.5 [Information Interfaces and Presentation]: Miscellaneous; J.3 [Life and Medical Sciences]: Medical Information Systems

## Keywords

visual analytics, information visualization, computer interpretable guidelines, guideline compliance, medical information systems

## 1. INTRODUCTION

Clinical Guidelines are developed on the basis of scientific evidence and aim to assist in the care process [24]. The application of guidelines in clinical practice, and their integration into information systems, that handle the patients' data, play a key role in order to increase the quality of health care and to support guideline design. Efforts have been made in the development of computer interpretable guidelines (CIG) [10] and in the visualization of patient data,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*i-Know '13* September 04 - 06 2013, Graz, Austria

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2300-0/13/09 ...\$15.00.

but approaches combining both are still rare (see section 2). Analysis of compliance with CIGs and the subsequent visualization of the results of this process remains an open and seldom considered issue. According to Quaglini [22, p.161], who defines compliance with a guideline as "acting according to the recommendations of a guideline", (non-)compliance plays an important role for the identification of possible reasons for implementation obstacles. Among these reasons there can be guideline quality, context, or user related issues. An important challenge is to make these huge amounts of available data accessible in an understandable visual way. The main contribution of this paper is an interactive visual exploration environment for the analysis of compliance with CIGs by providing:

- An action-based definition of compliance with three different types of actions.
- Visual abstraction of numerical values and highlighting of time points and intervals.
- Interval-based representation of repeatedly missing actions.
- Aggregated visual representation of compliance information.

At first we present the work related to our approach (section 2). Main design ideas of our approach are discussed in the beginning of section 3, followed by a short description of the example guideline, an outline of our action-based approach for compliance checking, and a detailed description of our solution for the visual representation of compliance information. We also discuss a case study conducted with a physician (section 4), and finally we provide a summary of our results and give an outlook to future work (section 5).

## 2. RELATED WORK

Our approach combines automatic data analysis (analysis of compliance with a CIG) with a visual representation of the analysis results within a unified view, in order to facilitate comparison of information about the patient, the guideline and CIG compliance. Various approaches exist for modeling CIGs, focusing on different tasks like guideline evaluation or decision support. Most of them use task networks models (TNM) and a formal definition of control flow and expression language [10]. We use a guideline formulated in Asbru, a plan representation language for the specification of computer interpretable guidelines [26, 28]. Asbru is intended to support tasks like guideline verification, validation, execution and critiquing. Asbru plans are composed of

time-oriented skeletal plans, which are meant to be instantiated for a treatment. It allows the modeling of preferences (plan selection), intentions (goals of treatment), conditions (model transitions e.g., from active to completed) and effects (relationship between plan or action execution and parameters). Plans can be nested inside other plans and can be executed in parallel or sequentially, and can contain cyclical (repeating) plan activations. Actions are atomic plans, they do not contain any subplans.

Various approaches for checking compliance with a guideline exist. Advani et al. [1] identify different levels of guideline based plan critiquing. Compliance can be checked, for instance, on the level of actions that have been taken in an actual treatment. This can be done by comparing executed actions and prescribed guideline actions with each other. This is the chosen approach for this work. On a higher level of interpretation the intentions of the guideline can also be taken into account (compared to the patient state, the actions, or the physician's intention) [1]. Different approaches related to ours exist for checking compliance by comparison of prescribed actions with actual executed ones. Chesani et al. [7] distinguish between events (i.e. actions) that should have happened, and events that should not have happened in the execution according to a formal guideline specification. Exceptions about these events are checked against an execution log and violations can be reported. Lucas et al. [13, pp.20-21] define a more fine-grained classification scheme for non-compliant actions, consisting of "actions unsupported by patient findings", "conflicting actions", "non-compliant order of actions", and "missing mandatory actions".

Visualizing the content of guidelines depends on the tasks and the users' groups [2]. Glare [31], for example, supports design and execution of guidelines by providing a flow chart-like visualization. AsbruView [17] shows temporal and hierarchical properties of Asbru guidelines with the help of 3D diagrams, and uses different metaphors for the plan components. Plan Strips focus on the hierarchical properties of plans by applying nested strips [25]. In our contribution we enrich a flow chart-like hierarchical plan structure view with compliance information, because this kind of visualization is especially suitable for showing aggregated information, and allows to give a quick overview (see section 3).

Visualization is commonly applied in the analysis of patient data, this data is often time-oriented and overwhelming in its amount. Dot plots are commonly applied for time-oriented data, notable examples are Graphical Summary of Patient Status [29] and CPDV [11] (and its successor MIVA[12]). We integrate compliance information in similar visualizations for patient data. Other approaches apply multiple coordinated views for quantitative as well as qualitative patient data. VisuExplore [23] is an example of how to effectively combine various visualization methods for time-oriented data, such as bar charts and line plots together with event and time line charts for qualitative data. Chittaro [8] presents a system especially tailored for mobile devices with means to aggregate data using different time granularities and to combine multiple data sources into one view. Temporal data abstraction poses an important issue when dealing with medical data, as reference to a priori knowledge is often required to interpret numerical data correctly. Data abstraction can be seen as a mapping from low level data to higher level concepts [30]. VIE-VENT uses temporal abstractions adapted to the context and derives high

level trend patterns [19]. Among systems that apply abstractions especially for visualization, are Knave-II [27] and Midgaard [4, 5]. Midgaard combines patient observation and treatment (plan) data, offers mechanism for semantic zooming, including data abstraction mapped to color in the lowest level of detail. Knave-II focuses on management and visual presentation of different abstractions for patient data. Abstractions classes are mapped to the vertical displacement of rectangles, whereas the horizontal length of the rectangles encodes the duration of time, during which the values fall into the same categorical class. We apply a similar visual abstraction technique for time-oriented numerical patient data to emphasize the guideline related context.

Compliance with a CIG and the data needed to check compliance are closely entangled with each other. Therefore we use highlighting to ease the identification of valid and invalid action time points and intervals, during which the application of an action is missing, inside patient parameter views. Applying highlighting and coloring in these manner can be considered as event visualization [9], including means for the detection and representation of events. EventViewer [6] is an approach using colored bars in order to mark time related events and allows rearrangement and organization of multiple coordinated views. Eventflow [20] visualizes aggregations of multiple event patterns with vertical bars (representing time points and intervals).

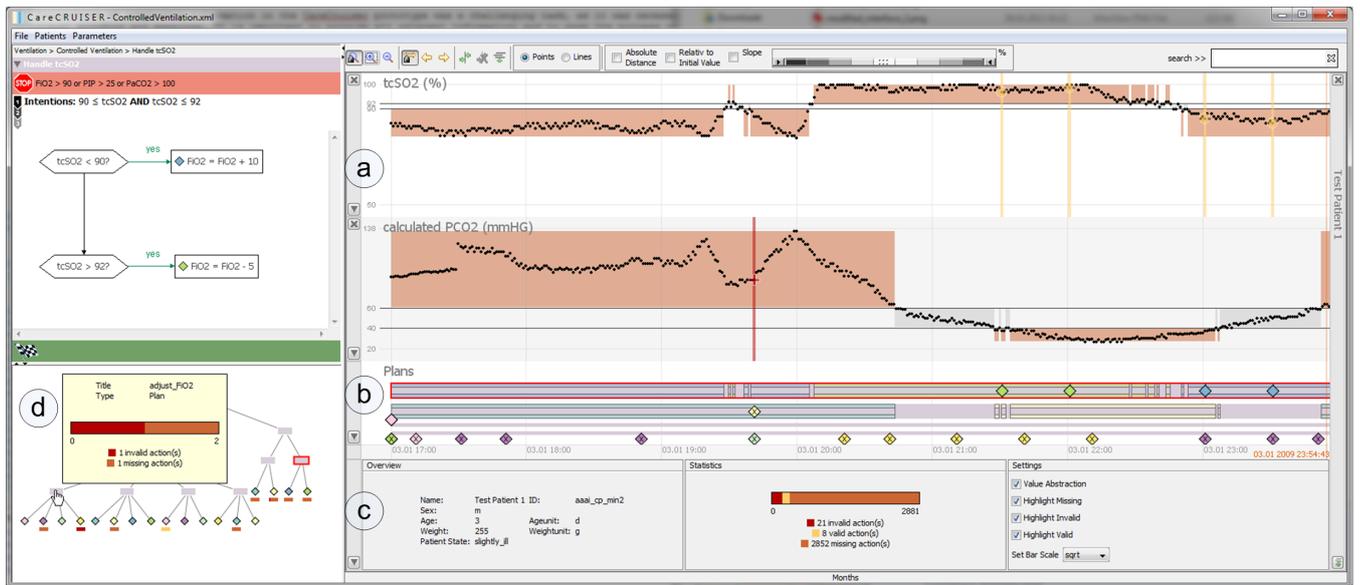
CareCruiser [14] aligns views for patient observations and actions executed during treatment on a common time-axis. It combines views for patient observations, treatment and CIG data. A major feature is the color-coded highlighting allowing to encode different information, i.e. distance to the intended value, progress from the initial value and slope between values. Furthermore, it provides views to inspect the structure and semantics of the related guideline. CareVis [3], the predecessor of CareCruiser, was the first tool to combine views for observation, treatment, and guideline data in this manner.

We selected CareCruiser as foundation for our extensions, because in our opinion it provides convenient and suitable combined views for patient data and CIGs. In the following section we outline our approach for the integration of compliance information into these views.

### 3. CONCEPTS OF DESIGN

Integration of compliance information in the CareCruiser prototype was a challenging task, the attention of the medical expert performing the analysis has to be focused on important aspects without hiding the underlying data. Information about compliance has to be integrated in a consistent manner, avoiding to overwhelm the user with too much information and aiming to help the user to identify important information quickly. We tried to solve this by different means (see Figure 1), considering the following aspects:

*Transparent overlay and highlighting in patient parameter views:* We apply abstraction of parameters into categorical classes (normal, low, and high), in order to visualize them as overlay in the raw patient data visualizations. This technique is meant to increase the speed of critical value identification, while preserving the information of the raw data view. The same is true for highlighting, emphasizing invalid/valid action time points and intervals, during which the execution of an action is missing. Highlighting can be activated or deactivated separately and aims to focus the



**Figure 1: System UI.** Value abstraction is integrated into patient parameter views (a). Highlighting of valid and invalid action time points and intervals of missing action application is supported. The modified plan execution view is placed beneath (b) invalid actions are marked accordingly (labeled with "X") and intervals of missing action execution are encoded by a bar consisting of an upper and lower part. A panel containing patient information, options and statistics is shown on the bottom (c). The statistics contain a stacked bar, showing the proportions between valid, invalid and missing action counts for the whole execution. Aggregation of valid, invalid and missing action counts is implemented in the overview (d). A stacked bar showing the aggregated counts for the current element (action or plan) is displayed inside a tooltip. Miniature versions of these bars are rendered directly beneath the actions for quick identification.

attention on important time points or time spans.

*Intervals of Missing Action Application:* The repeatedly executed plans for ventilation (containing actions applied with a short time delay) can generate a lot of individual missing actions; aggregating them into a missing interval helps avoiding visual clutter. Furthermore the exact time point of action application (between minimum and maximum delay) depends on user choice. Using intervals therefore also helps to cope with this indeterminism on conceptual level.

*Different types of actions:* Defining action types regarding CIG compliance (valid, invalid, or missing actions) allows to aggregate information. Details for a certain instance are omitted and one can concentrate on overall guideline compliance from a higher point of view.

*Higher level aggregation in the structural plan overview:* Aggregated compliance information for plans, subplans, and actions is integrated in the guideline structure view. It aims to give a quick overview for analysis tasks, which do not require the inspection of temporal relations between actions or parameters.

### 3.1 Example Guideline: Artificial Ventilation of Neonates

As an example, we refer to a guideline about artificial ventilation of newborn infants. We have used this specific guideline to compute compliance with a CIG and to visualize the results of this process. The overall intention of the guideline is to ensure the oxygen supply of the patient by applying a pressure gradient to the lungs. During this process the oxygen saturation ( $SO_2$ ) and the partial pressure

of carbon dioxide ( $PCO_2$ ) has to be kept within the normal range.

The guideline semantics can be separated into the initial phase and the phase of active ventilation. These phases are represented by two plans on top of the plan hierarchy, delegating specific task to their subplans. Each subplan of the initial plan is meant to initialize a certain parameter of the ventilation machine (depending on the patient state), related to the applied pressure (PIP, PEEP), to the fraction of oxygen in the air ( $FIO_2$ ) and the frequency of ventilation. The initial plans and the atomic clinical actions contained within them (initialization actions) have to be applied in sequential order. The controlled ventilation plan starts after the initial plan, and tries to stabilize  $SO_2$  and  $PCO_2$  values by applying two subplans. The plan handling the  $PCO_2$  parameter adjusts the ventilation frequency if the values are out of normal range (and this adjustments are repeated as long as the parameter does not return to normal range). The  $SO_2$  values are handled by a similar plan, which adapts the oxygen level ( $FIO_2$ ). The graphical representation of this plan is visible on the left of Figure 4. The abort condition on top (red) is fulfilled as soon as one of several parameters exceeds a critical value. The intentions beneath model the goals of the treatment (stabilize the  $SO_2$  parameter in a certain range). The if-then-else branch (conditions framed by hexagons) selects one of two mutually exclusive actions (framed by a rectangle). The first action (increase fraction of oxygen) is applied, if the oxygen saturation is below a certain value, the other action (decrease oxygen) is applied in the opposite case. Actions applied during treatment and

actions defined in plans have to be distinguished: Action instances are actual executed actions, an action defined in a CIG can therefore be instantiated multiple times during treatment. The plan is executed repeatedly according to a minimum and maximum delay time, as defined in the parent plan. The ventilation stops if either the conditions for plan completion or the conditions for aborting the plan get fulfilled.

## 3.2 Action Compliance with a CIG

Compliance in our contribution is checked on the level of (atomic) clinical actions, executed by a care giver (or by a machine). The actions prescribed by the guideline are checked against the user executed actions. Compliance checking represents the analytic component of our system, this is done on the basis of an imperative algorithm developed for the prototype. Contrary to usual data pre-processing in the field of information visualization (applies transformations based on raw source data), our approach also involves additional processing steps, requiring high level knowledge as represented by a CIG. This CIG data is used together with patient observation and treatment data to enrich the visualization according to the mantra of visual analytics "Analyze First—Show the important—Zoom, Filter and Analyze Further—Details on Demand" [16, p.82]. We define three types of different actions according to CIG compliance:

- *Valid Action*: The executed action is compliant with the guideline and applied correctly.
- *Invalid Action*: An action has been applied during the treatment, but should not have been applied according to the guideline. Executed actions with no equivalent in the guideline are automatically defined as invalid actions.
- *Missing Action/Missing Action Interval*: An action should have been applied, but is missing in the execution. In case of an interval, these are time-spans, during which an action should have been applied (repeatedly), but was not applied.

Computation of compliance is handled differently for the initial plans (containing sequential actions) and the ventilation plans (executed repeatedly).

*Initial plans* and their actions are applied sequentially, this means that they have to be applied in the correct order (if they are logged as applied at the same time point, it is assumed that they are applied in the correct order). An action is considered to be valid if it is executed in the right order and the action condition is true. An invalid action is an action that has been applied although its conditions have not been fulfilled (parameter propositions do not hold), or an action that has been executed, although later actions have already been applied. A missing action is generated if a later action is applied and the action has not been applied before.

The *ventilation plans* are repeated according to their minimum and maximum delay.

- An action is regarded as valid if the conditions for the action are fulfilled (parameter below the threshold or above the threshold) and if it is applied after reaching the minimum delay, i.e. if the temporal distance to the previously applied action of the plan is equal or greater than the minimum delay.

- An action is considered to be invalid, if their conditions do not hold, or if it is applied before the minimum delay is reached.

We define intervals of missing actions instead of individual missing actions (time points) for repeating ventilation plans. An interval representing the missing application of an action starts as soon as the following is fulfilled:

- The maximum delay has to be reached (the temporal distance to the previously applied action is greater than the maximum delay).
- The conditions for action application are true (but no valid action is applied)

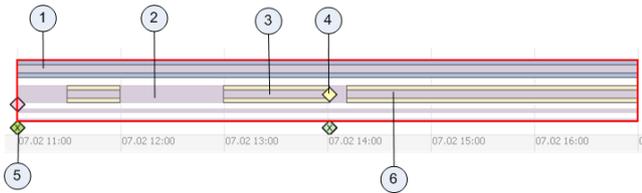
A missing interval is closed if the conditions are not fulfilled anymore, if the corresponding action is applied, or if the abort or complete condition gets fulfilled. The number of missing actions can be estimated from the intervals (e.g., dividing the interval time span by the maximum delay time).

## 3.3 Visual Encodings for Plan Execution

The treatment of a patient and the observation data recorded in course of it are often time-oriented. The plan execution view shows the actual executed actions (the treatment the patient has undergone) and time is mapped to the horizontal axis. We extended this temporal view to encode information about action compliance (see Figure 2). Valid actions (actions applied correctly during treatment according to the guideline) are represented by diamonds filled with color encoding the CIG action ID of the instance. Invalid actions (not applied accordingly) are marked with an "X", metaphorically stating that they are forbidden by the guideline. Time spans (intervals, during which the execution of an action is missing) are represented as two connected upper and lower parallel bars, leaving enough space for actions that might occur at the same time. The use of intervals instead of the generation of individual missing actions has several reasons. Repeating plans with a short delay time (e.g., a few seconds) can lead to a lot of missing actions (even for short time spans of missing action execution) and visual clutter. Aggregation into intervals avoids this issue. The other reason is that the exact time point of valid action execution is depending on user choice (the action can be applied anytime in the interval between the minimum and maximum delay). Individual missing actions generated for the visualization would only represent one possible solution. Furthermore we adapted tooltips to display additional information for the current element; counts for action instances having the same ID and compliance type (e.g., invalid) can be inspected this way.

### 3.3.1 Abort and Complete Conditions

The execution of the ventilation plans is halted as soon as complete or abort conditions are fulfilled (parameters are above or beneath a certain threshold). The semantics of abort and complete conditions are different, the former means that the plan failed, whereas the latter means the plan finished successfully. On the execution level the results are the same, as the plan(s) are halted as soon as the condition is true, for compliance checking this means that all subsequent actions are set to invalid and open missing action intervals are closed. We encode this transition to abort or completed state by altering the background color of a plan



**Figure 2: Plan execution view for an example patient.** The plan handling the oxygen saturation (1) and the plan handling the partial pressure of carbon dioxide (2) is shown. The action for adjusting low oxygen saturation is missing for the whole duration of the execution in the first plan (connected upper and lower bars in blue color, the plan is represented by the grey background). The other plan (2) has some time spans, during which the action for adjusting low values is missing (bars in yellow color). As soon as the valid action is applied (yellow diamond) (4), the missing interval ends (3). A new interval starts afterwards (6), because no valid action has been applied, although the parameter is still beneath normal range. Actions without an equivalent in the guideline (5) are automatically marked as invalid actions (diamonds marked with "X").

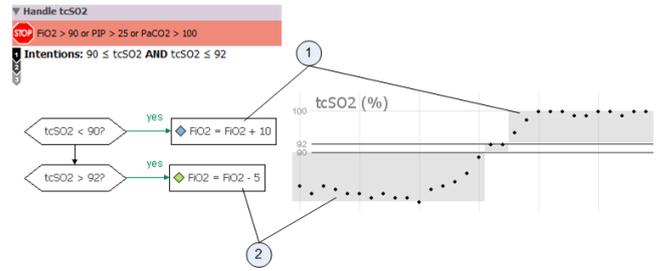


**Figure 3: Aborted plan.** The duration of halted execution, starting from time point of change to aborted state (blue bar representing a missing action interval ends) until the end of execution, is colored in red. All actions occurring in this time span are marked as invalid (diamond marked with "X"). Green color is used in a similar way for plan completion.

in the execution view. Red coloring for aborted state and green coloring for completed state are used to mark the time span of halted execution, from the beginning of the abort or complete state transition until the end of execution (see Figure 3). The colors used to mark the duration of abort (red) and complete plan (green) states correspond to the color of conditions in the logical detail view.

### 3.4 Parameter View Integration

Measured patient observations and applied clinical actions are usually closely related to each other. Analyzing compliance with a CIG therefore requires to integrate views for patient data, which is often numerical and updated frequently. We visualize the abstraction of numerical values into three categories (normal, high, low) as overlay (see Figure 4 on the right). The visual abstraction allows to identify immediately if a value is above or below normal range. The abstraction is not fully opaque and also allows seeing the parameter dot plot shining through. It is intended to support cognitive pattern recognition without hiding the details of the underlying data. Treatment actions of the two ventilation plans (monitoring the PCO<sub>2</sub> and SO<sub>2</sub> parameters) are applied as soon as values get out of range. As soon as values for PCO<sub>2</sub> or SO<sub>2</sub> fall into abstracted high or low categories, this goes hand in hand with the application of an action belonging to



**Figure 4: Abstraction and action conditions.** The parameter view with active overlay is shown on the right, the corresponding plan is shown on the left. Values of the SO<sub>2</sub> parameter mapped to the category high are associated with the application of a specific action (1) the same is true for values mapped to low category (2)

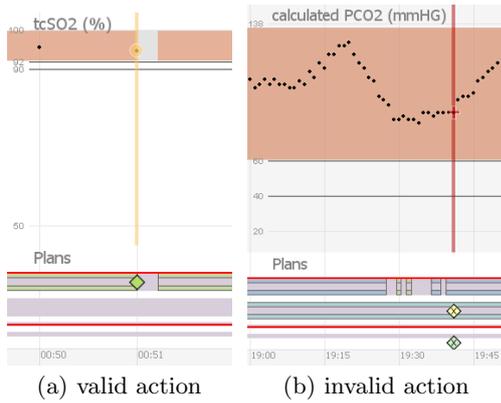
the corresponding plan (see Figure 4), this has the intention to assist in the detection of critical values and to help to see at one glance, if action execution is required.

Treatment, patient observations, and CIG compliance are closely coupled with each other. Highlighting of valid and invalid action time points is therefore supported in the associated parameter views: The analyst can directly see the consequences of valid or invalid action execution in the patient parameter view. This has been designed to make it easier to analyze the relationship between compliant/non-compliant action execution and parameter values. By applying this kind of highlighting, we try to prevent that experts have to switch between parameter and plan execution view too often during inspection. The type of action compliance (invalid or valid) is mapped to a transparent shape marking the parameter value (might be interpolated if no measurement is available for a specific time point) occurring at the highlighted action time point, and to the color. The purpose of the vertical bars is to help in the quick identification of executed action time points together with the parameter values related to them. Figure 5 shows an example execution in the parameter view, with highlighting activated for valid and invalid actions and missing action intervals.

Highlighting of missing action intervals has been designed to support the user in the identification of associated parameter values. The corresponding time span is highlighted inside the rectangle encoding the abstraction (see Figure 5 on the left), which means the occurrence of critical values together with absence of the appropriate actions is distinguishable from guideline compliant treatment during critical patient states. The progression of parameter values without the proper action applied can be examined inside the highlighted area and assist in analyzing how seriously the absence of the action influences the patient state.

### 3.5 Aggregation of Compliance Analysis Results

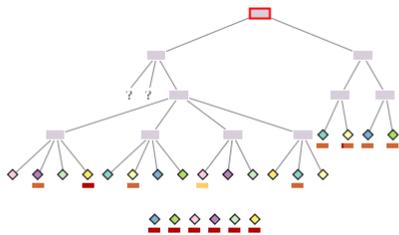
Specific analysis tasks which do not require further exploration of time dependent patient data are supported by additional visual cues. We provide an overview of valid, invalid and missing action counts (estimated from missing intervals) by mapping them to a stacked bar. Color coding is used for the compliance type using the same colors



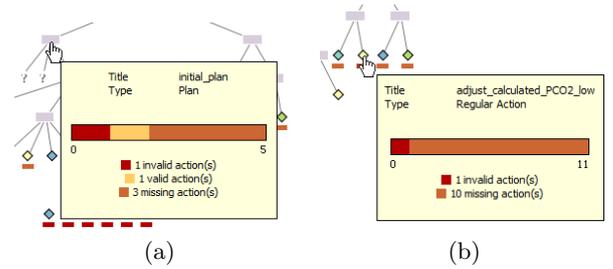
**Figure 5: Parameter view highlighting.** Highlighted valid action time point inside the SO<sub>2</sub> parameter view (yellow vertical bar, with circular shape marking the value at this time point) (a), these action causes the preceding missing action interval (light brown) to be closed. A new interval starts afterwards, because the action is not applied again, although the parameter is still out of range. An invalid action time point is highlighted in red (b) inside the PCO<sub>2</sub> parameter view (cross shape marking the value).

as for highlighting inside the parameter views. A horizontal stacked bar shows the proportion of counts for an action compliance type in relation to the total count. Without referring to quantitative information, one can quickly estimate the proportion of counts in relation to each other, in addition the exact counts are provided by a legend belonging to the bar. In order to avoid that high counts for a certain compliance type take up to much space and to keep the other counts visible, it is possible to use a square-root (instead of linear) scale for the counts. The overall counts for the whole execution are displayed in this way in a view on the bottom, additionally containing patient demographics and settings for highlighting and bar scaling (see Figure 1c).

### 3.5.1 Integration into Guideline Structure View



**Figure 6: Plan overview with (horizontal) miniature bars beneath the action symbols (diamonds).** It is clearly visible e.g., that a lot of guideline actions have been applied invalidly at least once (red bar) and that a lot of guideline actions are missing at least once (brown bar).



**Figure 7: Data aggregation in the guideline structure view.** A tooltip for the initial plan shows that most of the actions of the child plans are missing in this example (proportional area of stacked bar colored in brown) (a). Tooltip for the action, which is applied, if the PCO<sub>2</sub> Parameter is low, in this example the action is missing frequently and has one invalidly executed instance (red) (b).

CIG guidelines and treatment recommendations in general can consist of different sub units. The test guideline used in course of this project is formulated in the Asbru plan representation language, allowing guideline specification and decomposition into a set of skeletal plans, which can consist of sub plans. CareCruiser already provides a separate view, giving an overview of the guideline structure. It shows the hierarchical relationship of the plans contained in the guideline and the actions belonging to the guideline, as well as actions executed during treatment although they do not belong to the guideline. We used this structural view of the guideline to integrate information about aggregated valid, invalid and missing action counts for each plan and atomic actions. Compliance can be analyzed for each unit of the guideline by inspecting the aggregated counts. For an action this means that only valid, invalid or missing action instances with the corresponding action ID are counted, whereas for plans this means that the counts of all actions in the plan itself and for all actions contained in his subplans are aggregated together. Altered tooltips show an enlarged version of the stacked bar for the selected unit (see Figure 7). Miniature versions of these bars are rendered beneath each action (see Figure 6). They are intended to give a quick impression of the overall action compliance and to enable to see if an instance of an action e.g., has been applied invalid once in the execution.

## 4. CASE STUDY

We conducted a user case study with a medical expert (resident physician in neurosurgery), focused on gathering qualitative feedback and suggestions for future refinements. We were interested how an expert uses such a system, if he understands the visual metaphors represented in the system and is able to identify relevant patterns. Our approach for the case study oriented on evaluation methods used in usability testing [21], information visualization [18] and on the methodology applied for user test sessions with CareCruiser [15], in these sessions user tests were performed with domain experts and interviews have been conducted afterwards. The first stage of our methodology consists of a user test, relying on observation and the thinking aloud method (user is asked to continuously express his thoughts). The

interview is scheduled after the test, consisting of questions about the users impression, advantages and drawbacks in his opinion, and the request for comments and suggestions.

We first introduced the physician to the procedure, context, and purpose of our case study. Furthermore the guideline semantic was quickly covered and the physician was given the opportunity to explore the main interface and interactions of CareCruiser and the modified prototype for a short amount of time (to get an overview of the interaction mechanisms and the interface). We prepared a list of simple tasks for the physician, the tasks focus on evaluation goals related to information visualization systems, like the users comprehension and understanding of visual elements/metaphors. Two examples for this tasks are:

1. How often has the `adjust_tcSO2_high` action been applied validly? (Achieve this by using the plan execution view and locate the corresponding values in the parameter view.)
2. Find out, if the `adjust_tcSO2_low` action has been missing at least once, and/or has been applied validly at least once, and/or has been applied invalidly at least once.

The tasks were designed to cover all functionality and visual mappings introduced into the prototype. The physician managed to solve most of the tasks without problems. Difficulties occurred when he was asked to use the guideline structure view to seek aggregated counts for the initial plan: It seemed that he was not aware of the plan hierarchy in the guideline structure view and had difficulties identifying the correct plan. This might be avoided in future versions by making the meaning of the hierarchical elements more obvious (e.g., by labeling). Despite these issues, the miniature bars rendered beneath the action in the guideline structure view seemed to help to quickly identify if an action is missing. Solving tasks related to the temporal views (parameter and plan execution) seemed to work intuitively for the physician most of the time. Identifying valid actions in the temporal plan execution view and relating them with parameter values was quickly achieved by the test user. When he worked on identifying invalid guideline actions, he was not sure, which invalid actions belonged to the guideline and which were not in the guideline, this implies that he had problems with distinguishing them visually from guideline actions. Highlighting and encoding of missing intervals seemed to be understood intuitively, as he managed to accomplish related tasks very quickly.

The physicians' impression of the parameter and plan execution view modifications was very good, as he pointed out in the subsequent interview. Considering the highlighting he pointed out that "it is good to see the carried out actions and what they cause, to see the time of an event and the resulting consequences". He remarked that he would like labels in the plan execution view (e.g., inside a missing action interval) and for live systems he suggested hint boxes/messages containing cues like to increase a certain parameter. He stated that the visualization of compliance in the overview is generally useful, especially for offline analysis in his opinion. The interview as well as the usage observations within this case study showed that parameter view highlighting and the representation of invalid, valid, and missing action counts for the whole treatment seemed to be especially intuitive and useful for the physician.

## 5. CONCLUSION

We present a novel approach for integrating information about executed treatment compliance with a CIG (Computer Interpretable Guideline) into views for patient data and guideline structure of CareCruiser [14]. We combine analytical methods to determine compliance and visualization of the results. We suggest an action based definition of CIG compliance, as well as visual encodings and mappings for the generated information. Our approach was driven by the following main design ideas: abstraction of high frequency parameter data, highlighting of important time points and intervals, and aggregation of data and action counts to provide a quick overview.

A user case study was conducted with a physician. The presented prototype was generally appreciated; suggestions and found issues give input for further improvements. It turned out that most of our visual encodings are useful and help to identify relevant patterns quickly. Adaptation of the plan overview, like labeling, might help to emphasize the meaning of the elements in future versions. In addition tighter visual separation will be considered between actions that do not appear in the guideline and actions that are included in the guideline, but are incorrectly applied. Moreover, possible adaptations of interface color schemes will be investigated. Representative user studies gathering more data need be conducted in the future to gain further insights. Future work in general might focus on the visualization of different definitions of compliance or on finding additional means for the visualization of guideline adherence on various levels of abstraction.

## 6. ACKNOWLEDGMENTS

This work was carried out as part of the MobiGuide project funded in part by the European Commission under the 7th Framework Program, grant #287811. We would like to thank Dr. med. Daniel Firouzabadi, resident physician in neurosurgery, for his participation in the case study and his suggestions and ideas. Furthermore we would like to thank Theresia Geschwandtner for providing the source code of CareCruiser and Albert Amor-Amorós for his feedback and comments.

## 7. REFERENCES

- [1] A. Advani, K. Lo, and Y. Shahar. Intention-based critiquing of guideline-oriented medical care: The asgaard project at stanford. In *Proc. AMIA Annual Symposium*, pages 483–487, 1998.
- [2] W. Aigner, K. Kaiser, and S. Miksch. *Visualization Techniques to Support Authoring, Execution, and Maintenance of Clinical Guidelines*, page 140–159. IOS Press, Health Technology and Informatics, 2008.
- [3] W. Aigner and S. Miksch. Carevis: integrated visualization of computerized protocols and temporal patient data. *Artificial intelligence in medicine*, 37(3):203–18, 2006.
- [4] W. Aigner, A. Rind, and S. Hoffmann. Comparative evaluation of an interactive time-series visualization that combines quantitative data with qualitative abstractions. In *Computer Graphics Forum*, volume 31, pages 995–1004. Wiley Online Library, 2012.

- [5] R. Bader, S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '04*, pages 105–112, New York, NY, USA, 2004. ACM.
- [6] K. Beard, H. Deese, J. Emerson, and N. Pettigrew. The eventviewer: A tool for visualizing and exploring events extracted from ocean observing system data. In *OCEANS 2010 IEEE - Sydney*, pages 1–8, 2010.
- [7] F. Chesani, E. Lamma, P. Mello, M. Montali, S. Storari, P. Baldazzi, and M. Manfredi. Compliance checking of cancer-screening careflows: an approach based on computational logic. *Studies in health technology and informatics*, 139:183–92, 2008.
- [8] L. Chittaro. Visualization of patient data at different temporal granularities on mobile devices. In *Proceedings of the working conference on Advanced visual interfaces, AVI '06*, pages 484–487, New York, NY, USA, 2006. ACM.
- [9] T. Christian. Event-based concepts for user-driven visualization. *Information Visualization*, pages 65–81, 2011.
- [10] P. D. Clercq, K. Kaiser, and A. Hasman. *Computer-Interpretable Guideline formalisms.*, volume 139 of *Studies in Health Technology and Informatics*, pages 22–43. IOS Press, 2008.
- [11] A. Faiola and S. Hillier. Multivariate relational visualization of complex clinical datasets in a critical care setting: A data visualization interactive prototype. In *Tenth International Conference on Information Visualization, 2006. IV 2006.*, pages 460–468, 2006.
- [12] A. Faiola and C. Newlon. Advancing critical care in the icu: a human-centered biomedical data visualization systems. In *Proceedings of the 2011th international conference on Ergonomics and health aspects of work with computers, EHAWC'11*, pages 119–128, Berlin, Heidelberg, 2011. Springer-Verlag.
- [13] P. Groot, A. Hommersom, P. J. F. Lucas, R.-J. Merk, A. ten Teije, F. van Harmelen, and R. Serban. Using model checking for critiquing based on clinical guidelines. *Artif. Intell. Med.*, 46(1):19–36, May 2009.
- [14] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. Carecruiser: Exploring and visualizing plans, events, and effects interactively. In *Pacific Visualization Symposium (PacificVis), 2011 IEEE*, pages 43–50, 2011.
- [15] 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 *Proc. of the BCS HCI Conference*, 2011.
- [16] D. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual data mining. chapter Visual Analytics: Scope and Challenges, pages 76–90. Springer-Verlag, Berlin, Heidelberg, 2008.
- [17] R. Kosara, S. Miksch, Y. Shahar, and P. D. Johnson. Asbruvew: Capturing complex, time-oriented plans - beyond flow-charts. In *The Second Workshop on Thinking with Diagrams 1998 (TwD-98)*, pages 119–126. The University of Wales, The University of Wales, 1998.
- [18] R. Mazza. *Introduction to Information Visualization*. Springer Publishing Company, Incorporated, 1 edition, 2009.
- [19] S. Miksch, W. Horn, C. Popow, and F. Paky. Context-sensitive and expectation-guided temporal abstraction of high-frequency data. In *Proceedings of the Tenth International Workshop for Qualitative Reasoning (QR-96), Stanford Sierra Camp. Fallen Leaf Lake, CA*, pages 154–63. AAAI, 1996.
- [20] M. Monroe, K. Wongsuphasawat, C. Plaisant, B. Shneiderman, J. Millstein, and S. Gold. Exploring point and interval event patterns: Display methods and interactive visual query, 2012.
- [21] J. Nielsen. *Usability Engineering*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.
- [22] S. Quaglini. *Compliance with clinical practice guidelines.*, volume 139 of *Studies in Health Technology and Informatics*, pages 160–79. IOS Press, 2008.
- [23] A. Rind, W. Aigner, S. Miksch, S. Wiltner, M. Pohl, T. Turic, and F. Drexler. Visual exploration of time-oriented patient data for chronic diseases: Design study and evaluation. In *Proceedings of USAB 2011: Information Quality in e-Health*, pages 301–320. Springer, Springer, 2011.
- [24] K. Rosenbrand, J. V. Croonenborg, and J. Wittenberg. *Guideline development.*, volume 139 of *Studies in Health Technology and Informatics*, pages 3–21. IOS Press, 2008.
- [25] A. Seyfang, K. Kaiser, T. Gschwandtner, and S. Miksch. Visualizing complex process hierarchies during the modeling process. In *Business Process Management Workshops*, pages 768–779. Springer, 2013.
- [26] A. Seyfang, R. Kosara, and S. Miksch. Asbru's reference manual, asbru version 7.3. Technical report, 2002.
- [27] Y. Shahar, D. Goren-Bar, D. Boaz, and G. Tahanl. Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artif. Intell. Med.*, 38(2):115–135, Oct. 2006.
- [28] Y. Shahar, S. Miksch, and P. Johnson. The asgaard project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. In *Artificial intelligence in medicine*, pages 29–51, 1998.
- [29] S.M.Powsner and E. R. Tufte. Graphical summary of patient status. *Lancet*, 344(3):386–9, 1994.
- [30] M. Stacey and C. McGregor. Temporal abstraction in intelligent clinical data analysis: A survey. *Artif. Intell. Med.*, 39(1):1–24, Jan. 2007.
- [31] 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*, 101:162–6, 2004.