

TimeSpan: Using Visualization to Explore Temporal Multi-dimensional Data of Stroke Patients

Mona Hosseinkhani Loorak, Charles Perin, Noreen Kamal, Michael Hill, and Sheelagh Carpendale

Abstract— We present TimeSpan, an exploratory visualization tool designed to gain a better understanding of the temporal aspects of the stroke treatment process. Working with stroke experts, we seek to provide a tool to help improve outcomes for stroke victims. Time is of critical importance in the treatment of acute ischemic stroke patients. Every minute that the artery stays blocked, an estimated 1.9 million neurons and 12 km of myelinated axons are destroyed. Consequently, there is a critical need for efficiency of stroke treatment processes. Optimizing time to treatment requires a deep understanding of interval times. Stroke health care professionals must analyze the impact of procedures, events, and patient attributes on time—ultimately, to save lives and improve quality of life after stroke. First, we interviewed eight domain experts, and closely collaborated with two of them to inform the design of TimeSpan. We classify the analytical tasks which a visualization tool should support and extract design goals from the interviews and field observations. Based on these tasks and the understanding gained from the collaboration, we designed TimeSpan, a web-based tool for exploring multi-dimensional and temporal stroke data. We describe how TimeSpan incorporates factors from stacked bar graphs, line charts, histograms, and a matrix visualization to create an interactive hybrid view of temporal data. From feedback collected from domain experts in a focus group session, we reflect on the lessons we learned from abstracting the tasks and iteratively designing TimeSpan.

Index Terms—Multi-dimensional data, Temporal event sequences, Electronic health records

1 INTRODUCTION

Working closely with domain experts, we have designed, implemented and studied TimeSpan, to develop a better understanding of temporal data of acute stroke patients. Stroke is the second leading cause of death globally and the major cause of acquired neurological disability in adults [17]. Fast and efficient treatment of stroke patients can reduce stroke related mortality and disability. Ischemic stroke is caused by a sudden blockage of a brain artery. It can be treated with tissue plasminogen activator (tPA), but this is a time critical treatment. Rapid administration of tPA to open a blocked artery in the brain will be beneficial on average within 4.5 hours of stroke onset. However, treatment must commence as soon as possible and all delays from onset-to-hospital and hospital arrival-to-treatment must be minimized.

The time from when a patient arrives in the hospital to when tPA is administered is called door-to-needle (DTN) time. There are various delays in DTN time due to patient and hospital related factors [8]. Examples of such delays include delay in obtaining CT scan, delay in patient registration, delay due to patient high blood pressure, and delay in getting blood lab results. Multiple small delays that a patient encounters may add up to a large delay in onset-to-treatment.

We are working with a group of stroke professionals who are studying clinically acquired temporal stroke treatment data to better understand the varying time spans in DTN. Understanding the factors that contribute to these delays needs careful examination and analysis of the temporal multivariate data. Improving the support for exploration of these data may contribute to finding unforeseen reasons for delay, and lead to novel approaches to providing faster treatment. Our mandate is to design a tool that can help this professional group with their analysis.

Currently, the standard technique for representing and analyzing these data are statistical process control (SPC) charts [4] borrowed from industry. SPC charts aggregate data to give an overview of the

data and make it possible to perform statistical tests on this data. Using this method, many hospitals have improved the quality of care [30]. However, SPC charts do not demonstrate the detailed timing of events, cannot represent every patient in detail, and do not include the multi-dimensional aspects of the data. Crucial information can be missed or not observable because of aggregation. Indeed, stroke patient data is multi-typed, consisting of temporal, ordinal, quantitative, and nominal data types, making it a challenging problem in data visualization.

To design TimeSpan, we conducted a series of observations in the emergency department of a large tertiary-care hospital and eight one-on-one interviews with stroke professionals with various expertise (*e. g.*, stroke neurologist, quality assurance analyst, and stroke nurse). These studies provided an understanding of the current practices and challenges of our target domain. From the interviews, we extracted, analyzed, and classified the analytical tasks a visualization should support. Based on this classification and iterative feedback, we extracted design goals that informed the design of TimeSpan to support health professionals in their data exploration tasks. The purpose of TimeSpan is to support exploratory analysis of the temporal and multi-dimensional data of stroke patients, outside of clinical hours so that the stroke team can change and improve the stroke treatment process.

To assess the benefits of TimeSpan, we conducted a focus group with five stroke specialists. We discuss the focus group results in terms of design decisions and suitability for exploration of the data. We observed that working with TimeSpan generated considerable excitement about the tool as well as a new understanding of the power of visualization. It also triggered the domain experts ask new questions and resulted in an expanded requirements and tasks list, including new factors triggered by exploring the visualization. Visualization acted as a catalyst by initiating new ideas that the stroke group had not envisioned before.

The three main contributions of this paper are:

1. A list and classification of basic tasks in the critical domain of stroke care analysis that can be used by subsequent researchers to visualize data of stroke patients, and that might apply to similar problems, *e. g.*, door to balloon data of heart attack patients.
2. TimeSpan, designed in close collaboration with domain experts, combining multiple visualizations to support exploration of multi-dimensional, multivariate, and temporal data of stroke patients.
3. A set of lessons we learned during a focus group session introducing TimeSpan to the stroke team.

2 METHODOLOGY

We were approached by the stroke team because they were interested in discovering whether new data analysis tools might be useful to them

• Mona Hosseinkhani Loorak, Charles Perin, and Sheelagh Carpendale are with the Department of Computer Science, University of Calgary. E-mail: hossem, charles.perin, sheelagh@ucalgary.ca.

• Noreen Kamal and Michael Hill are with the Department of Clinical Neurosciences, University of Calgary. E-mail: nrkamal, michael.hill@ucalgary.ca.

Manuscript received 31 Mar. 2015; accepted 1 Aug. 2015; date of publication xx Aug. 2015; date of current version 25 Oct. 2015.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.




Fig. 1. Simplified process flow diagram for treating stroke patients.

in the search for new ways to improve stroke patients' outcome. Thus, we started our collaboration to explore how visualization of stroke treatment data might be helpful. Our target audience are the stroke team members including stroke neurologists, quality assurance specialists, triage nurses, and Emergency Medical Service (EMS) members.

During our initial investigation, we found that the approach to stroke care was highly team-based. Our target experts perform different tasks to improve the quality of care (e.g., process changes, sending out weekly performance reports to the team, and regular focus group discussions). To assess their current practices and discover the tasks they need to perform to analyze their data, we conducted eight interviews with stroke experts from the hospital. Three of them are stroke neurologists responsible for recognizing stroke patients and administering the right treatment; one is a postdoctoral fellow; two are quality assurance analysts working with the stroke team to make the process of treating stroke patients more time-efficient; one is a unit manager of the emergency department who helps to build a process to deliver tPA to patients as soon as possible; and one is a triage nurse responsible for observing strategy improvement processes, collecting patient data, and presenting the data to the stroke team.

We also performed a series of observations and contextual inquiries in the emergency department of the same hospital. Furthermore, we participated in two of their focus group meetings, in which they were discussing the positive and negative points in their recent cases and how things could have been improved for individual patients.

Based on these interviews and observations, we first understood the process for treating stroke patients, familiarized ourselves with the dataset, and analyzed the limitations of current practices for analyzing their data (section 3). We reviewed related work in visualization and concluded that no tool is well-suited to the analysis of multi-dimensional stroke patient data (section 4). Then, we classified the tasks and requirements (section 5) to inform the design of TimeSpan (section 6). Finally, we conducted a focus group session with stroke team members, that was designed to be similar to an envisioned use case, to assess the benefits and limitations of TimeSpan (section 8).

3 CURRENT PROCESS AND DATA PRESENTATION

In this section, we briefly describe the process of treating acute stroke patients, present our dataset, and describe the current method of presenting this dataset to the members of the stroke team.

3.1 Process for Treating Stroke Patients

There are several steps involved from the onset of stroke symptoms to when the patient is treated by either administering tPA or endovascular therapy (also known as groin puncture) or both of them. Figure 1 shows a simplified diagram of the process flow for treating stroke patients.

tPA is most effective within 4.5 hours from the time of onset, yet there is declining benefit as time elapses. The standard DTN time is set by the Brain Attack Coalition to be 60 minutes [1]. However, recent research shows that among 641 US hospitals, only 6.7% treated more than half of their patients within 60 minutes [12]. The delays are either patient-related or system-related. Examples of patient delays include patients having high blood pressure that needs to be managed before administering tPA, or patient being on anticoagulants, requiring the laboratory testing to be confirmed prior to determining eligibility to receive tPA. System-related delays include delays in diagnosis, in obtaining a CT scan, and in communication and action.

To improve the time to treatment, the stroke team periodically alters the process of delivering tPA in order to gradually reduce treatment time intervals. The effects of these changes need to be analyzed in order to determine if they should be integrated into the process permanently. Examples include sending a pre-notification by EMS services about

Weekly Summary for each patient




Fig. 2. A sample PDF file showing the current visual representation of stroke patient data sent out weekly to the stroke team.

the patient and the estimated arrival time to the emergency department (STATStroke), delivering tPA in the CT scan room, and registering the patient as unknown at arrival time to speed up the admission process.

3.2 Dataset

The dataset consists of 150 stroke patients between June 2013 and July 2014 who all received intravenous tPA as part of their treatment. Some of these patients also underwent endovascular treatment in addition to tPA. For each patient, two types of information are recorded:

1. **Temporal attributes:** Include temporal data points in minutes such as arrival date at the hospital, last time the patient has been seen as normal (LSN), time of calling the hospital by EMS services regarding the arrival of a patient (EMS dispatch time), door time, CT scan time, tPA time, and endovascular treatment time.
2. **Multi-dimensional Patient attributes:** Consist of three types of data: 1) quantitative data such as age and stroke severity (NIHSS) ranging from 0-42; 2) ordinal data such as the level of consciousness (LOC), speech, and facial status; and 3) nominal data such as the location of receiving tPA, arrival mode, and whether the patient was registered as unknown.

Some patients are associated with additional unstructured textual information regarding their status and with the reasons why their treatment was delayed. Because the data for each patient are gathered manually by the stroke team during the time the patient arrived and went under treatment, some attributes are missing in the dataset.

3.3 Current method for Presenting Data

Our interviews revealed that domain experts believe that increasing people's awareness of their performance improves their overall outcomes, e.g., "we implemented just reporting about the times in a tabular format about a year and half ago. And around that time, we saw about a 10 minute decrease in our average treatment times" (P08). Currently, stroke team members receive weekly reports regarding their performance both for the last week and the last eight months. These reports used to be in the form of tabular views. However, the team recently created a static visual report by abstracting some detailed data and representing aggregated data using line charts and bar charts (Figure 2).

Five of the experts found these visual reports to be useful, in particular to quickly see if they are meeting their goals, e.g., "I think it's great because this shows the DTN time that we are trying to achieve and shows easily how far we are over that goal." (P02). Two experts would like to have access to more detailed data but were concerned that the representation might get overwhelmed.

Two experts preferred the tabular view of the data, observing that visual reports aggregate the data while there is a need to access the details of each attribute, e.g., "I personally like this one better [pointing at the tabular data representation]. Just because this one gives you the specific details, [...] this one [pointing at the tabular view], here is your delay, patient was intubated and waiting for INR. That sort of thing. This one [pointing at the visual report] gives you high level numbers but it does not give you the reasons why." (P05)

Finally, one stroke expert preferred a hybrid of the tabular view and the visual report to get both an overview and the detail of each patient: *“I like a kind of hybrid of both. I mean this one [pointing to the visual report] is really nice because it shows you the context of past performance as well as current performance. [...] I might like a hybrid, because the table to me is fairly clear in terms of what is going on. I can look at it and see the patient details line by line” (P08).*

4 RELATED WORK

Although the literature in visualization has mostly not addressed the visualization of stroke patient data, there has been a growing body of work in the context of health [3, 10, 13–15, 19, 22, 23, 33]. Rind et al. provide an extensive review of visualization tools for health data that they group into two categories [26]: those that focus on individual patients (*e. g.*, LifeLines [22], graphical summary of patient status [23], MIVA [11], and VisuExplore [25]), and those that focus on a group of patients. Given our problem, we discuss the second category. These visualizations focus on visualizing data from a group of patients in order to assist clinical researchers and quality assurance analysts in assessing and improving the quality of care of patients [13, 15, 19, 33, 36].

PatternFinder [10] is an interface for specifying temporal queries with value and time span constraints. Queries are specified through form-based interfaces and the results are in the form of ball-and-chain visualization (node-link diagrams with horizontally aligned nodes and thick edges showing temporal data). PatternFinder shows the detailed information of each patient matching the query. Since this system focuses on queries, it does not provide an overview of patients’ data.

Lifelines2 [33] represents each patient record on a horizontal time line based on the time that the event occurred. It also provides several interaction techniques such as temporal summaries, alignment, and filtering for data exploration. This system focuses on representing the temporal events with timestamps, not on the temporal intervals between two timestamps. Moreover, events in Lifelines2 can occur several times for a given patient, which is not the case for stroke patient data.

LifeFlow [36] and its successor EventFlow [19] extend Lifelines2 by aggregating data from patients whose corresponding sequence of events are similar. The Aggregation view highlights the common patterns that appear in event sequences. EventFlow extends LifeFlow by encoding the average time of events for the aggregated patients using vertical bar charts. Both of these systems provide overview+detail of the dataset. They are designed to handle varying sequences of events, while in the case of stroke patient data, event sequences are determined and ordered.

OutFlow [35] and its successor DecisionFlow [13] adopt Sankey Diagrams [24] to analyse disease progress and outcomes in patient records. DecisionFlow extends OutFlow by using aggregation techniques for the visual encoding of high-dimensional temporal event sequence data. Both tools aggregate symptoms and their average development time for the patients in color-coded paths. Each path shows the outcome of each symptom with statistical information within the patients’ population. Both systems aggregate patients at each step of the disease by similarity and are designed to handle varying sequences of events.

VISITORS [15] is based on KNAVE [27] and KNAVE II [28]. It uses an ontology-based exploration and aggregation module to visualize patients’ raw data and abstract concepts from them. The purpose of VISITORS is to assist in quality assessment of clinical trials by exploring derived concepts and their associations in patient data. The evaluation of the system with domain experts highlights the necessity of having a simpler visualization with complex mode capabilities.

While these related works have made significant contributions to the visualization of cohort patient records, none of them focus on fluctuating time intervals within a stable sequence of events. Only Patternfinder and Lifelines2 can visualize the detailed timing of individual patients, and most of the existing systems aggregate patient data in order to discover patterns in event sequences. As shown in Figure 1, stroke patient data follows a pre-determined and ordered sequence of events.

As no precise requirement and task description exists for stroke patient data, we extracted requirements and classified the tasks that stroke experts must perform in order to inform the design of TimeSpan (section 5). As part of our requirement analysis, we discovered that

experts require a visual representation of both an overview and detailed information of patients—the detailed information being heterogeneous multivariate attributes. However, to the best of our knowledge, visual encoding of multi-dimensional attributes of patients together with their temporal data has not been well-explored in the literature.

5 TASK ANALYSIS

From the transcripts of interviews and observations, we extracted a series of analytical tasks that a visualization tool should support to help the stroke professionals improve patients’ quality of care.

From the gathering of 40 tasks, we extracted and classified a set of 26 representative tasks, that are presented in Figure 3. The column “data” indicates whether the data is available or not. We identify three types of data that are involved in performing the tasks (columns in “data type”): 1) date or time (*e. g.*, date of patient arrival at the hospital and month of February 2014); 2) temporal attributes (*e. g.*, DTN time and CT to Needle time); and 3) patient attributes (*e. g.*, age and arrival mode). We used Bertifier [21] to group the tasks by data type. We obtain five groups of tasks: 1) tasks involving patient attributes only (T1–4, blue); 2) tasks involving temporal attributes only (T9–14, yellow); 3) tasks involving dates only (T20–26, red); 4) tasks involving both patient attributes and temporal attributes (T9–14, green); and 5) tasks involving both temporal attributes and dates (T15–19, brown).

We use Amar et al.’s taxonomy of low-level analytic tasks [2] to describe each (often compound) task into fundamental analytic tasks (columns in “corresponding low-level task”). Each representative task involves a small number of low-level tasks, but each low-level task is required at least once. The most frequent low-level tasks are Filter, Sort, Correlate, and Characterize. The least frequent ones are Determine range, Find extremum, and Cluster. This correspondence between representative and low-level tasks informs the design of interactions. The columns “corresponding interaction” are Yi et al.’s interaction categories [37]. Several interaction categories could be used to perform some tasks, and we indicate which interaction category we implemented in TimeSpan. This classification informed the design of our prototype and provides a guide for future research in similar application areas.

6 THE TIMESPAN SYSTEM

We describe TimeSpan with its design goals that we derived from our observations, interviews and iterative discussions with domain experts.

- DG1 Keeping familiarity:** The visualization components should be familiar and easily understood by various health and medical professionals involved in analysis of acute care of stroke patients.
- DG2 Complexity on demand:** While keeping the visual data representations familiar is one of our main goals, we need ways to represent the multi-dimensional and multi-typed data. Thus, our goal is to increase the representation complexity on demand.
- DG3 Integrating heterogeneous data:** The data attributes are heterogeneous. Ways of representing numerical, nominal, and ordered data vary significantly and our goal is to integrate these different data types altogether.
- DG4 One holistic view:** Patient and hospital related attributes affect the temporal data recorded for each patient. Thus, the goal is to design a holistic view embedding both the temporal data representation and the multi-typed data attributes.
- DG5 Access to the raw data:** Health professionals are trained at rapidly comprehending large amounts of text. Thus, the raw data should always be accessible, in textual or numerical format.

We designed TimeSpan to fulfill the identified design goals and support the classified analytical tasks. The stroke team problem—discovering the factors that are introducing delays in the treatment process—is by nature exploratory. Therefore, our rationale is to represent all the data, and all the temporal, multivariate and multi-typed dimensions in a single holistic view, assuming that every single factor matters in exploring and understanding the data. We made our design decisions based on iterative paper prototyping of various alternatives and asking for experts’ feedback to compare different possibilities.

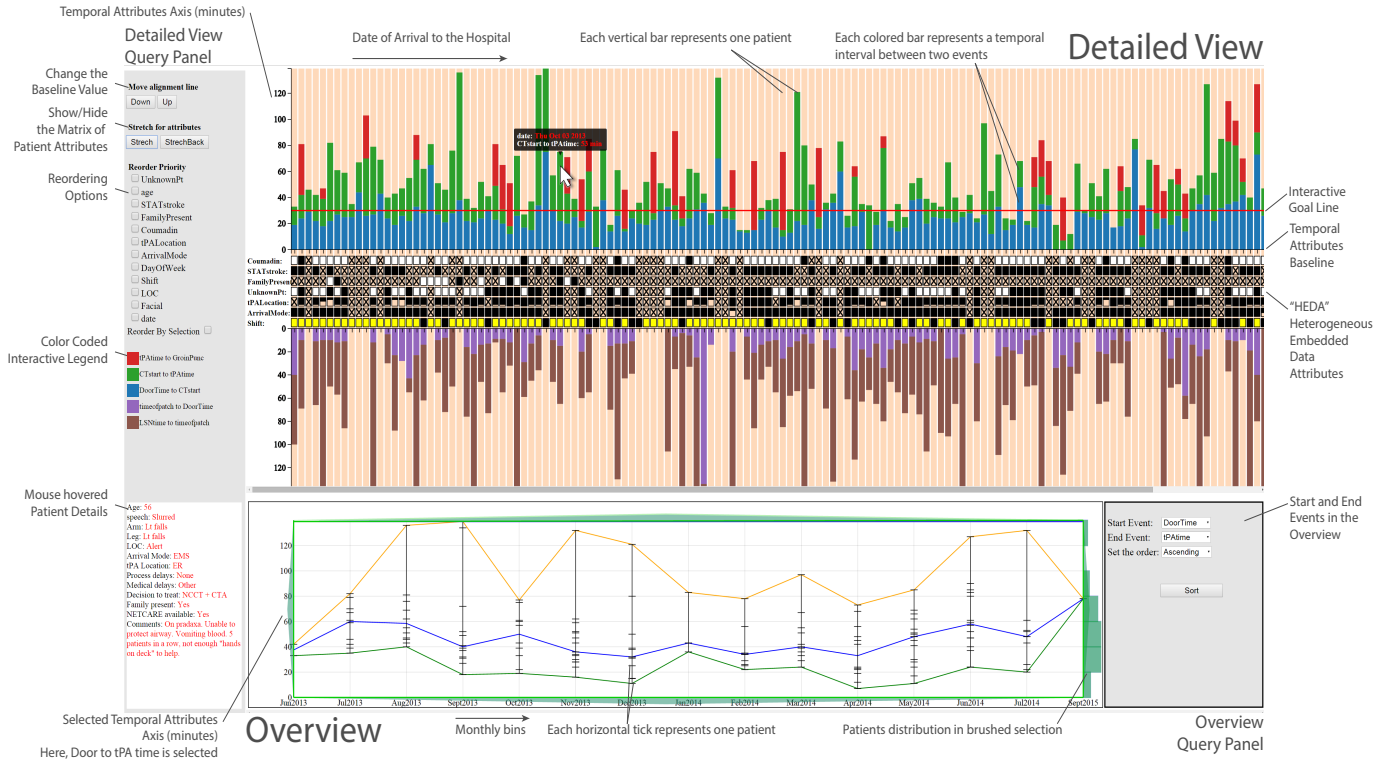


Fig. 4. Overview of TimeSpan, consisting of the detailed view and the overview, with their associated query panels.




Fig. 5. Brushing and linking two selections (S_1 and S_2) in the overview updates the patient backgrounds in the detailed view with the same color as the selection tool. Non-selected patients are faded out.

a more aggregated view (see Figure 2). Thus, the overview shows monthly bins on the horizontal axis (T22).

The vertical axis encodes time, in minutes. By default, the vertical axis represents DTN time. However, the ability to explore different event time intervals is significant for domain experts (T19). Changing the values of start and end times in the overview query panel updates the vertical axis to this new selected time range. Each horizontal tick within a month represents a patient, giving both a count (T23) and an overview of the patients distribution within each month (T24, T26). The minimum, maximum, and the median time of patients (according to the selected vertical time range) in different months are shown using line charts with green, orange and blue colors respectively (T15, T16).

Several tasks involve filtering patients according to a time period (T15, T16, T23–26), a temporal event duration (T17), or both (T18). Brushing and linking selection tools are available to perform these tasks. The selection tools are rectangular areas whose width and height specify both a date range and a temporal event time range. As T19 and T24 involve comparing two time periods, TimeSpan features two

selections that can overlap. Patients within a selection are linked in the detailed view by matching their background color to the color of the selection. Figure 5 shows two selections, where patients within the blue selection S_1 have a blue background, patients within the green selection S_2 a green one, patients within both selections ($S_1 \cap S_2$) an orange one, and patients who are not within either of the two selections ($\neg(S_1 \cup S_2)$) have a white background and are faded out.

Selections are also useful to provide alternative strategies for performing some tasks. For example, T11 can be performed by setting the upper border of a selection to 60 minutes, in which case patients with a time range below this value will be selected and other patients will be faded away. Finally, in order to make it possible to count and compare the distribution of two subsets of patients (T23, T25, T26), we enrich the selection areas with histograms on their right hand side border.

6.3 Additional Interactions

TimeSpan features two additional powerful interactions making the tasks easier and faster to perform: reordering and the interactive legend.

Hierarchical reordering to reorder the columns of the HEDA. Reordering rows and columns in tabular data representations can aid in revealing patterns [21]. The list of reorderable attributes (see Figure 4, *detailed view query panel*) can be reordered by dragging and dropping attribute labels. Thus, patients can be hierarchically reordered with the position of the attributes being used to set attributes priority. For example, one could select *age* and *STATstroke* and move *STATstroke* above *age* to first reorder patients based on *STATstroke*; and in each group of values for *STATstroke*, reorder patients based on their *age*. Reordering makes it easier to perform tasks such as T20 and T21 (sort operations), T1–8 (reconfigure), and T10–13 (explore hierarchical reordering).

Interactive legend in the query panel showing the color associated with each temporal interval. Fading out temporal intervals that are not relevant for a given task is useful for exploring some specific temporal event periods such as door to needle or CT scan to groin puncture time.

7 TIMESPAN WALK-THROUGH

To illustrate how TimeSpan can be used for visual exploration of stroke patient data, we describe a walkthrough scenario based on real-world stroke patient data. The scenario we use parallels the most frequently




Fig. 6. Setting the baseline to be tPA administration time.




Fig. 7. A green bar crossing the red goal line means that the corresponding patient had a CT to tPA time longer than 40 minutes.




Fig. 8. Comparing distributions of patients DTN time in 2013 and 2014.

described scenario from our experts. A quality improvement expert, Sarah, is responsible for reviewing the data through this tool approximately every one or two weeks to see if changes that they are making to the process are creating improvements. She would use the tool to discover the underlying factors that might cause lower or higher DTN times. She would then feed this back to the entire team to alter their Continuous Improvement Strategy. Sarah opens the TimeSpan web page where all the patient data are visually represented and updated weekly. By a quick look at the detailed view, she recognizes that some colored bars are unexpectedly longer than usual (T10,T12) and wants to access detailed information for the corresponding patients. Placing her mouse on top of a colored bar, she reads the information in a tooltip and in the textbox on the query panel (T3,T14). She notices that the details for these longer bars vary considerably, triggering further exploration.

Sarah then explores the tPA to groin puncture times (red bars) because groin puncture occurs sporadically. She presses the “Down” button twice to push the bars lower placing the baseline at tPA time (Figure 6). She discovers that between January and April 2014, the patients either did not go through endovascular treatment or the tPA to endovascular time was shorter than usual (T13). She then decides to find out which patients had a CT scan to tPA longer than 40 minutes (T11). She presses the “Up” button to set the baseline at CT scan. She drags the red goal line to 40 minutes. Since she is only interested in the green bars (CT scan to tPA), she unselects the tPA to groin puncture time in the interactive legend. The red bars fade out. She can see, by noticing whether or not the green bar crosses the red goal line, that only fifteen patients had CT to tPA time longer than 40 minutes (Figure 7).

Now Sarah wants to compare the DTN time of patients in 2013 and 2014, to see if the treatment process has improved (T19). She uses the blue selection in the overview to select patients in 2013 and the green one to select patients in 2014 (T25,T26). She compares the number of patient distributions in the blue and green selections by looking at the histograms on the side of each selection (T23,T24). She finds that in 2013 the largest number of patients had a DTN time between 20 to 40 minutes, while in 2014 more patients had a DTN time between 40 and 60 minutes (Figure 8), suggesting that the DTN time did not improve.

Intrigued by this finding, she decides to narrow her search comparing the month of July 2013 (blue selection) with the month July 2014 (green selection) in terms of DTN time (T25,T26,T19). The detailed view is updated to reflect these selections by changing the background color of each month. Patients from July 2013 are on the left side of the detailed view, and those from 2014 are on the right side, making comparison difficult. Thus, she brings the two selections together in




Fig. 9. Selecting and reordering by selection patients from July 2013 (blue selection) and 2014 (green selection) to compare their DTN time.




Fig. 10. Reordering patients based on selection (only patients from 2014 are selected), then based on shift.




Fig. 11. Patients who got tPA in the CT scan room (full black cells) have less CT to tPA time compared to the rest of the patients.

the detailed view by clicking the “Reorder by selection” checkbox in the detailed view query panel. Selected patients are adjacent, grouped by the selected month (T22), and non-selected patients are faded out on the right hand side (Figure 9). From this exploration, Sarah discovers that CT scan to tPA time of patients in July 2014 is generally shorter than patients in July 2013 (T17). However, tPA to groin punctures are longer and more frequent in July 2014 than July 2013.

Moving on, she decides to explore the patient attributes looking for potential patterns and exploring how the patient attributes affect the treatment times (T6). She clicks on the “Stretch” button in the detailed view query panel to make the HEDA appear. Sarah wonders if patients in 2014 were admitted more frequently during the day or during the night. She first uses the green selection tool to select patients who were admitted in 2014, and presses the “Reorder By Selection”. Then, she reorders the patients according to whether they arrive during the day or night by checking the “Shift” checkbox. This reorders the patients according to shift within the selection (T1), with first the day (yellow cells), then the night (black cells). She observes that approximately two thirds of the patients in 2014 arrived during the day (Figure 10).

Additionally, Sarah wonders if CT scan to tPA time is related to tPA location (T7). She keeps the same selection and reorders the patients based on tPA location. Then, she fades out the tPA to groin puncture time to focus on CT scan to tPA time only. Sarah discovers that patients who got tPA in the CT scan room (full black cells) have less CT to tPA time compared to the rest of the patients (Figure 11).

Sarah thinks that another factor might be involved as there are some exceptions. She reorders patients by “STATstroke” (pre-notification to the hospital) as well (T5). She sets the reordering priority to be “tPALo-




Fig. 12. Patients who had STATstroke (black cells) and received their tPA in the CT scan room (black cells) are having better CT to tPA time compared to the patients who received tPA in CT but did not have STATstroke (crossed cells).



Fig. 13. Selecting patients with CT to tPA time lower than 40 minutes and sorting them by CT to tPA time (green bars). Long CT to tPA times seem to involve shorter and less frequent endovascular therapy (red bars).

cation”, then “STATstroke”, then “shift”, then “date” by moving the reordering attributes in the detailed view query panel (Figure 12). She finds that the patients who had STATstroke (black cells) and received their tPA in the CT scan room have better CT to tPA time compared to the patients who received tPA in CT but did not have STATstroke.

Towards the end of the exploration, Sarah changes the data represented in the overview. She changes the start and the end events using the two drop-down lists in the overview query panel to be CT time and tPA time, respectively. She wants to focus on patients who had a CT scan to tPA time shorter than 40 minutes. She sets the green selection to horizontally select all patients, and to vertically select only patients with a CT to tPA time lower than 40 minutes (T18). The patients who are filtered out are faded out in the detailed view. Finally, she sorts the patients based on their CT to tPA times in an ascending order using the “Sort” button. Through this exploration, she finds out that patients who have longer CT scan to tPA times tend to have relatively smaller tPA to groin puncture times, and they have less frequent endovascular therapy compared to the rest of the patients (Figure 13).

8 FOCUS GROUP RESPONSE

In order to receive feedback about the benefits and limitations of TimeSpan, we presented it to stroke experts in a two hour focus group session.

8.1 Methodology

Of the different possible ways of eliciting responses we chose a focus group [18] because focus groups offer an informal environment where participants exchange ideas. Also, since one of the envisioned uses of this visualization is for the stroke team to examine the data in order to discuss the process; a focus group emulates this usage. Focus groups offer a combination of gathering expert opinions and observing expert brainstorming. Focus groups are also known to tend towards group-think [9, 32], where discussions often converge to mutually acceptable ideas and outlying ideas may not be expressed. In our situation, since

we are looking for new design directions, such convergence of mutually acceptable ideas is potentially useful. We are aware that care of stroke patients is a team work and that of necessity patient care is always primary. Thus, assessing TimeSpan with a team process seemed appropriate. Of the five experts, three (two stroke neurologists and a post-doctoral fellow) took part in the initial interviews and the two others (a stroke neurologist and a stroke nurse) were experts we were meeting for the first time.

During the focus group, we explained the visualization using a sample use case scenario based on the tasks we had collected. We videotaped the session and transcribed the video. Then, we analyzed the transcription grouping comments according to common ideas and concepts as in affinity diagrams [6]. The experts paid close attention to the explanations and asked on the spot questions, and partly because of their deep familiarity with the data, they rapidly found TimeSpan understandable. As a result, not much time was needed to develop an understanding of the different mappings and interactions. Instead the discussion concentrated on the visualization system’s potential.

8.2 Results

In this section we present the results of the focus group session¹.

8.2.1 Overall Response

The general tone of the focus group was extremely positive. The experts described several different ways in which they saw this tool being used, e. g., “This is a very good tool for visualizing our own data and hospital process and every hospital should have this kind of data. Should be standardized.”, “I could envision kind of a weekly or a monthly meeting where we would review cases and the system would show what the data looked like ...”, and “We could use the system for audit purposes. Why did this patient or this individual take so long ...”. Since this project’s main intention was to help this team of stroke experts in their efforts to improve the care of stroke patients, this quote effectively sums it up, “... This is a great tool which creates an easy way of visualizing data which is informative for us in driving change at the hospital”.

8.2.2 The Need for New Data

When the experts realized that the cross signs in the HEDA show that a given data item was not available, they noted that missing data was wide spread. “It seems that we have a lot of missing data”. Importantly, they were aware of the implications of this missing data. Since they currently use aggregation techniques on their data, seeing the attributes in the HEDA was revealing. They realized that they may be drawing conclusions upon small subsets of data, e. g., “If there are things that are mostly just missing data, I don’t know, you know maybe that’s not a useful data or we will just have to start collecting it right”. However, they were also aware that collecting data, where patients’ needs must come first was an ongoing challenge. The experts also discussed improving data consistency that might enable them to distinguish different factors such as in-hospital stroke patients from the rest of the patients: “The other thing that would be interesting, is in-hospital strokes. So people that have strokes in the hospital. And so, you have an onset time and a treatment time but you do not have a door time”. During the discussion, they indicated that stroke severity (NIHSS) is one of the most significant attributes they would like to explore further. However, they realized that while this meta-data was intended to be collected, it was not recorded in the dataset. These revelations were possible due to the HEDA providing detailed non-aggregated patient attributes.

Recurrently during the focus group session, experts noted the lack of some information that would require collecting new data. For instance, one expert was looking for higher granularity in time event intervals: “Is it also possible to consider even for example, in a big hospital area you can consider onset time, 911 call, dispatch time, scene arrival time, scene departure time, arrival to the hospital time.”. Another expert mentioned that patient blood pressure and how it has been handled could be visually encoded within TimeSpan: “I think blood pressure is

¹More detailed description of our observations can be found at <http://innovis.cpsc.ucalgary.ca/supplemental/timeSpan/>.

important to put on there or treatment for blood pressure”. As another instance, one expert was interested in collecting data about patients who have more than one recorded door time: *“It would be interesting to show patients with 2 door times who are transferred between hospitals. So, there is arrival at the first center and then arrival at the comprehensive stroke center. Then, we can actually see how we are doing as a system”*.

During the early interviews, we learned that the process of mixing and administrating tPA is not the same during the day and night. Therefore, we represented the patient arrival time at the hospital as either “Day” or “Night”. The experts found this dimension interesting, and discussed how this data could be of interest and how it might provide a better understanding of changes in process: *“[...] say 7 to 5pm or 8 to 5pm. So that’s a regular shift. Off hours you can actually break it up even more than that. You can say sort of 5pm to midnight and then your midnight to 8am. Cause I think you know 5 to midnight might still be fast but when everybody is asleep in bed there might be delays associated with that. And then weekend as well. Cause weekend is not regular working hours.”* The stroke team was interested in this type of data, but encoded in a manner that revealed the changes in process.

8.2.3 Insights Leading to New Questions

During the focus group, the experts’ exploration of the data led them to raise many new questions based on their increased understanding of their collected data. These new insights led quickly to ideas for further investigation. For example, during our field observation and interviews, we were told that the process of delivering tPA always followed this sequence: symptom onset, EMS dispatch time, door time, CT scan, tPA, and endovascular therapy. Using TimeSpan, one could see that the timings were not always in accordance with the above sequence of events. After our experts looked through the visualization and visually recognized a particular patient, they discussed about how the treatment process could sometimes change and that those cases are interesting for them. *“The ordering may not always happen such that tPA will happen before groin puncture. Sometimes tPA is given in the angio suite after groin puncture as well.”* Another expert said *“In the visualization, it would just have to overlap the red and the green like there is one there.”* This type of variation such as the possibility of order change in the events is just the type of process revelation that would allow the experts to examine whether the different ordering had an impact on treatment.

In showing the experts how to use the HEDA to explore the data they were interested to see that by sorting the patients based on their age with the baseline set at tPA time, the frequency of the red bars –representing tPA to endovascular time– are higher on the left hand side (younger patients) compared to the right hand side (older patients). This shows that most of the patients who got endovascular therapy were among the younger patients and that senior patients were less likely to receive it. After discussing this finding, the experts explained that stroke neurologists might already know about this fact: *“older people have more torturous arteries. So, likely they would have less endovascular treatments than younger patients. I think typically you tend to be more aggressive with younger patients”*. However, one stroke neurologist explained that this fact is worth further investigation as there might be some age-bias: *“everybody has certain pre-conceived notions on why we are doing things ... if they all had torturous vessels and they cannot get in there maybe yeah that is the explanation but maybe we are just age biased and we are slower with older patients”*. Another expert found it insightful and suggested that they need deeper analysis of the data. *“We would have to probably investigate more deeply ... did they truly have torturous vessels or not.”*

8.2.4 New Questions Leading to New Tasks

Inspired by the interactions with the system, experts also suggested new tasks. For example, one of the stroke neurologists mentioned that she wants to know how she has performed in terms of treatment time: *“I want to know how I’m doing”*. She also said that she is interested in comparing her performance to other neurologists: *“I want to have my median time compared to others”*. These comments were followed by a long discussion between the experts about the pros and cons of visualizing individual professional results and comparing them. As

another example, some patients with longer dispatch to door times caught our experts’ attention because they are having longer transfer times compared to the rest of the patients. An expert found a very long purple bar showing a long transfer between the dispatching time to when the patient arrives at the door of the hospital. Hovering over the purple bar showed that the length is equal to 134 minutes. *“The transfer times are generally really short so I’m guessing all of those are [Hospital X]. But look at this guy here, look at this transport time. You know that he is coming from [Hospital Y] or somewhere. There is no way that was from [Hospital X]”*. Afterwards, the experts started thinking about specifying and identifying the category distance (e. g., zone1, zone2, and zone3 within the city) and having some radius zones be collected, integrated, and visualized with TimeSpan.

9 DISCUSSION

Results from the focus group session emphasize several important points triggered by exploring the data with TimeSpan. While there were many factors that were initially surprising to the experts (large amounts of missing data, changes in process sequence, 2 door times, impact of time of day and age of patient), on discussing these points, because of their familiarity with the data, the experts could understand how the collection happened – in some case remembering an actual patient. However, the important point is that these insights raised new questions, which in turn, when investigated, may lead to changes in process. This response is normal in medical care situations where changes must be considered carefully because of potential impact on patients. Thus, new leads on what to investigate are extremely positive results. We discuss how visualizing data in a holistic way generally helped the domain experts become more aware of the importance of data quality and of the power of data. They also noted how accessing these data can reveal individual performances and be helpful or problematic. Then, we discuss results specific to the design of TimeSpan, informing the design of future similar systems, before providing future directions.

9.1 Data Empowerment

Throughout the focus group session, visualizing the data using TimeSpan made experts increasingly aware of the power of their data. This led them to consistently ask for new data, more data, and more complete data. They realized that they have questions to ask that needed more data than they had collected to date. Thus, they started thinking about the new data that could be gathered and included, all the while considering the difficulties inherent in collecting medical data.

Data quality awareness through visualization: Currently, the experts are analyzing their data and improving their performance awareness using tools, which aggregate data. However, with TimeSpan they noticed that these tools only revealed some aspects of the data, neither revealing how sparse the collected data is nor the relationships between multi-dimensional variables. Through TimeSpan’s holistic view of the temporal and multi-dimensional attributes in TimeSpan they became aware of both the quality and sparseness of their data and discussed improving their data collection: *“have to start collecting it right.”*

Visualization revealing the power of data: The experts noticed that the visualization can reveal the potential power of their current data. This in turn was driving requests for more tasks and visualization capabilities. The experts discussed how TimeSpan might help the stroke team members be more aware of their own performances, of the overall quality of care improvement, and help them better understand what works and what does not work in their current process, e. g., *“The most powerful thing you can do when we are doing quality improvement is to tell people: look this is what is actually happening. Cause people have false conceptions about how well they are doing. They do. And that is normal.”* Overall, we found that showing the data to the domain experts made them more aware of the power of data visualization: *“It’s amazing there’s so much variance, there’s nothing like looking at the data. I mean you could make some hypotheses”*.

Performance assessment: During the focus group, TimeSpan triggered questions among the experts about their own performance and about how they are doing compared to their colleagues. Some wished to know more regarding their performance: *“I want to have my median*

time compared to others . . . Median is important and you can compare and see if it's really bad." However, other experts pointed out the risks of interpreting such data, e.g., "The problem is you can see this huge barrier, there's so many contributing factors to delay that you can't really put it down to physician.", "So, if you get all the horrible patients from that month that just happen to come in [...] that doesn't mean that you are slower than the next guy." The experts discussed the power of data visualization in conjunction with the risks of being evaluated based on this data—how it can make people aware of how they are doing, but that all factors are not included—as a side effect of trying to improve the stroke team process.

9.2 Factors of Success—Lessons Learned

TimeSpan was quickly and unanimously accepted by the experts. The experts rapidly passed the visualization understanding barrier and most of the focus group session was dedicated to discussing the data, getting insights, and generating new question and tasks. Here, we discuss the factors that we think contributed to the success and acceptance of TimeSpan, as these can inform the design of future visualization tools dedicated to domain experts.

Starting with simplicity: The initial view consists of details in bar charts and an overview in line charts, thus the exploration starts with simple, familiar visualizations (DG1). The experts easily understood all the components and interactions of TimeSpan, including the overview and the detailed view, the visual mappings, and the interactive baseline. Starting from familiar charts made it easy to initially understand the visualization and immediately start explorations.

Simplicity to complexity: In our system description, we moved our explanation from simple to complex interactions. Due to its visual complexity, the HEDA was the last part demonstrated.

Since the experts were deeply interested in the interrelationships between all aspects of their data, one of our design goals was to include all the different types of heterogeneous data, while keeping comparative interaction viable (DG3). Although the various types of data have different representation techniques in the literature, but we showed them all in a unified way. As expected, experts had initial difficulties understanding the HEDA, e.g., "... I find it really confusing, the X's and the darkness and it doesn't seem completely intuitive to me to sort of unpack what all that means". However, while initially disconcerting, it quickly was perceived as the strongest feature of the system. This is the feature that generated the most discussions among the experts and, once it was fully understood, the experts were quite enthusiastic "... if you say – I am not interested in that one or that one – you can simplify it ... I like that. I like it a lot".

The lesson here is that while the complexity of the HEDA was daunting initially, the experts were readily capable of both understanding it and seeing its potential. Being able to gradually increase visual complexity did lead to empowerment as well as expert satisfaction.

Simplicity to complexity to simplicity: One of the main observations from the focus group is that experts continuously went back and forth between simplicity (bar charts and line charts) and complexity (multivariate attributes in the HEDA) (DG2). Basically, by looking at the simple graphs they could generate questions. To answer these questions, they then looked at the HEDA, formulated queries, and re-ordered the patients. Once the query was completed, they went back to the simple graphs to read the answer to their question.

This decoupling between the question/answer reading visualizations and the query specification visualization is possible because the patient attributes are seamlessly integrated into the bar charts and behave as an extension of the bar charts themselves. This is not the case, for example, when using multiple views where different dimensions are spatially separated. Thus, in TimeSpan the correspondence between a patient temporal events and the same patient attributes is made straightforward.

This seamless integration of all the data dimensions into one holistic view (DG4) anchored the familiar and relatively unfamiliar visualizations such that multi-dimensional patient attributes are shown in the same context as the temporal ones. The lesson is that integration of visualization techniques may cause complexity but can also empower data exploration when reading through familiar visualization techniques.

9.3 Future Directions

A frequent topic of conversation in the focus group was about data completeness: about missing data, incomplete data, and additional data they would need for their new questions. Real world data is usually messy and incomplete. The data we were provided, while not as complete as the experts would have liked, was of sufficient quality for us to work with. Moreover, the data is being collected manually by people directly involved with patient care. These front line care givers must of necessity put patient care first. In fact, considering the urgent conditions under which this data must be collected, the sample they gave us is remarkable. Thinking about possible technological help for the data gathering process is an interesting avenue for future research.

Because the experts quickly understood TimeSpan, they started to think as visualization designers and provided us with potential improvements for the system based on the new data, new questions, and new tasks they came up with. Such future improvements include: accessing some aggregated numerical data about the selected patients using the selection tools; being able to map the absolute arrival time on the x axis of the detailed view as an alternative to the current relative arrival time on the x axis; ability to create separators in the detailed view to focus on a contiguous subset of patients; being able to aggregate temporal events, e.g., merging door to CT scan to tPA time into a single event door to tPA time; and incorporating spatial dimensions (e.g., distance zones) with the current temporal approach.

Our experts also discussed deployment possibilities and strategies for TimeSpan as part of their working process. They envisioned having the tool running on one of their machines in the hospital for weekly group discussions. They appreciated that TimeSpan is a web-based tool requiring no particular expertise to run.

10 CONCLUSION AND FUTURE WORK

This paper introduced TimeSpan, a visual analysis tool designed and implemented in close collaboration with stroke experts to better support them in exploring and analyzing the temporal, multi-dimensional, and multi-typed data of stroke patients. We designed the system based on a set of observations in a local hospital and one-on-one interviews with domain specialists. From the analysis of these studies, we extracted a set of tasks and requirements that a visualization system should support to help domain experts in their decision-making process—whose deep understanding of their data may increase the quality of care of stroke patients. Using a combination of stacked bar graph and Heterogeneous Embedded Data Attributes, we visually depicted the temporal events of stroke patients and their multi-type attributes in a single unified view.

We reported the results of a focus group session with domain experts and reflected the lessons we learned during this collaboration. We received enthusiastic feedback from members of the stroke team, who actively participated in discovering insights about their data and discussing the benefits that TimeSpan will bring into their system. As a result of this session, experts realized the power of data and the importance of its quality. By looking through the visualization, they started thinking about collecting new data and more complete data. Moreover, based on the visualization, the experts came up with new realizations about their data, new questions, and new tasks. We also discussed the design lessons we learned while designing and evaluating TimeSpan. In Summary, the visualization lessons are 1) start with simplicity, 2) move to complex visual structures on demand, and 3) allow people to go back and forth between simplicity and complexity on demand.

Based on the enthusiasm of the stroke team members and their discussion about deploying the tool in the hospital, we hope that TimeSpan will be quickly adopted by stroke experts for exploring and analyzing the temporal and multivariate data of patients. It would be an interesting future work to adapt TimeSpan for use in other scenarios where individuals go through a set of events in a process and analysts want to explore the temporal as well as the multi-dimensional aspects of data, such as door to balloon data of heart attack patients.

ACKNOWLEDGMENTS

This research was supported in part by AITF, NSERC, GRAND, Surfnet, and SMART Technologies.

REFERENCES

- [1] M. J. Alberts, R. E. Latchaw, and A. Jagoda. Revised and updated recommendations for the establishment of primary stroke centers a summary statement from the brain attack coalition. *Stroke*, 42(9):2651–2665, 2011.
- [2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, INFOVIS 2005*, pages 111–117, Oct 2005.
- [3] R. Bade, S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. In *Proc. CHI '04*, pages 105–112. ACM, 2004.
- [4] J. C. Benneyan, R. C. Lloyd, and P. E. Plsek. Statistical process control as a tool for research and healthcare improvement. *Quality and Safety in Health Care*, 12(6):458–464, 2003.
- [5] J. Bertin. Readings in information visualization. chapter Graphics and Graphic Information Processing, pages 62–65. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [6] H. Beyer and K. Holtzblatt. *Contextual Design*. Morgan Kaufmann Publisher, Inc., San Francisco, 1998.
- [7] B. Buxton. *Sketching user experiences: getting the design right and the right design: getting the design right and the right design*. Morgan Kaufmann, 2010.
- [8] J. Desai and E. Smith. Prenotification and other factors involved in rapid tpa administration. *Current Atherosclerosis Reports*, 15(7), 2013.
- [9] J. K. Esser. Alive and well after 25 years: A review of group-think research. *Organizational Behavior and Human Decision Processes*, 73(2–3):116 – 141, 1998.
- [10] J. Fails, A. Karlson, L. Shahamat, and B. Shneiderman. A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. In *IEEE Symposium On Visual Analytics Science And Technology*, pages 167–174, Oct 2006.
- [11] A. Faiola and C. Newlon. Advancing critical care in the icu: A human-centered biomedical data visualization systems. In M. Robertson, editor, *Ergonomics and Health Aspects of Work with Computers*, volume 6779 of *Lecture Notes in Computer Science*, pages 119–128. Springer Berlin Heidelberg, 2011.
- [12] G. C. Fonarow, E. E. Smith, J. L. Saver, M. J. Reeves, D. L. Bhatt, M. V. Grau-Sepulveda, D. M. Olson, A. F. Hernandez, E. D. Peterson, and L. H. Schwamm. Timeliness of tissue-type plasminogen activator therapy in acute ischemic stroke patient characteristics, hospital factors, and outcomes associated with door-to-needle times within 60 minutes. *Circulation*, 123(7):750–758, 2011.
- [13] D. Gotz and H. Stavropoulos. Decisionflow: Visual analytics for high-dimensional temporal event sequence data. *IEEE TVCG*, 20(12):1783–1792, Dec 2014.
- [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)*, pages 43–50. IEEE, March 2011.
- [15] D. Klimov, Y. Shahar, and M. Taieb-Maimon. Intelligent visualization and exploration of time-oriented data of multiple patients. *Artificial Intelligence in Medicine*, 49(1):11 – 31, 2010.
- [16] S. Lin, J. Fortuna, C. Kulkarni, M. Stone, and J. Heer. Selecting semantically-resonant colors for data visualization. In *Proc. EuroVis '13*, pages 401–410. Eurographics, 2013.
- [17] R. Lozano, M. Naghavi, and K. Foreman. Global and regional mortality from 235 causes of death for 20 age groups in 1990 and 2010: a systematic analysis for the global burden of disease study 2010. *The Lancet*, 380(9859):2095 – 2128, 2013.
- [18] C. Marshall and G. Rossman. *Designing Qualitative Research*. SAGE Publications, 1999.
- [19] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman. Temporal event sequence simplification. *IEEE TVCG*, 19(12):2227–2236, Dec 2013.
- [20] T. Munzner. Marks and channels. In *Visualization Analysis and Design*, chapter 5, pages 94–114. CRC Press, 2014.
- [21] C. Perin, P. Dragicevic, and J.-D. Fekete. Revisiting bertin matrices: New interactions for crafting tabular visualizations. *IEEE TVCG*, 20(12):2082–2091, Dec 2014.
- [22] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, and B. Shneiderman. Lifelines: using visualization to enhance navigation and analysis of patient records. *Proc. AMIA Symposium*, pages 76–80, 1998.
- [23] S. Powsner and E. Tufte. Graphical summary of patient status. *The Lancet*, 344(8919):386 – 389, 1994. Originally published as Volume 2, Issue 8919.
- [24] P. Riehmann, M. Hanfler, and B. Froehlich. Interactive sankey diagrams. In *IEEE Symposium on Information Visualization, INFOVIS 2005*, pages 233–240, Oct 2005.
- [25] 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 A. Holzinger and K.-M. Simonic, editors, *Information Quality in e-Health*, volume 7058 of *Lecture Notes in Computer Science*, pages 301–320. Springer Berlin Heidelberg, 2011.
- [26] A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman. Interactive information visualization to explore and query electronic health records. *Foundations and Trends in HCI*, 5(3):207–298, 2011.
- [27] Y. Shahar and C. Cheng. Intelligent visualization and exploration of time-oriented clinical data. In *Proc. HICSS-32, 32nd Annual Hawaii International Conference on Systems Sciences*, volume Track4, pages 12 pp.–, Jan 1999.
- [28] Y. Shahar, D. Goren-Bar, D. Boaz, and G. Tahan. Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artificial Intelligence in Medicine*, 38(2):115 – 135, 2006. Temporal Representation and Reasoning in Medicine.
- [29] B. Shneiderman. The eyes have it: a task by data type taxonomy for information visualizations. In *Proc. IEEE Symposium on Visual Languages*, pages 336–343. IEEE, Sep 1996.
- [30] J. Thor, J. Lundberg, J. Ask, J. Olsson, C. Carli, K. P. Härenstam, and M. Brommels. Application of statistical process control in healthcare improvement: systematic review. *Quality and Safety in Health Care*, 16(5):387–399, 2007.
- [31] E. R. Tufte and P. Graves-Morris. *The visual display of quantitative information*, volume 2. Graphics press Cheshire, CT, 1983.
- [32] M. E. Turner and A. R. Pratkanis. Twenty-five years of group-think theory and research: Lessons from the evaluation of a theory. *Organizational Behavior and Human Decision Processes*, 73(2–3):105 – 115, 1998.
- [33] T. Wang, C. Plaisant, B. Shneiderman, N. Spring, D. Roseman, G. Marchand, V. Mukherjee, and M. Smith. Temporal summaries: Supporting temporal categorical searching, aggregation and comparison. *IEEE TVCG*, 15(6):1049–1056, Nov 2009.
- [34] T. D. Wang, C. Plaisant, A. J. Quinn, R. Stanchak, S. Murphy, and B. Shneiderman. Aligning temporal data by sentinel events: Discovering patterns in electronic health records. In *Proc. CHI '08*, pages 457–466. ACM, 2008.
- [35] K. Wongsuphasawat and D. Gotz. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE TVCG*, 18(12):2659–2668, Dec 2012.
- [36] K. Wongsuphasawat, J. A. Guerra Gómez, C. Plaisant, T. D. Wang, M. Taieb-Maimon, and B. Shneiderman. LifeFlow: Visualizing an overview of event sequences. In *Proc. CHI '11*, pages 1747–1756. ACM, 2011.
- [37] J. S. Yi, Y. a. Kang, J. Stasko, and J. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE TVCG*, 13(6):1224–1231, Nov. 2007.