

Echo: A Large Display Interactive Visualization of ICU Data for Effective Care HandOffs

Manu Mathew Thomas*
Department of Computer
Science
University of Illinois at
Chicago

Thomas Kannampallil†
Department of Family
Medicine
University of Illinois at
Chicago

Joanna Abraham‡
Department of Biomedical
and Health Information
Sciences
University of Illinois at
Chicago

G. Elisabeta Marai§
Department of Computer
Science
University of Illinois at
Chicago




Figure 1: (Left) Two collocated testers establish a common information space using Echo on a tiled large display wall that runs SAGE2 middleware [11]. (Right) A group of testers run Echo synchronously on a large display screen, on a laptop and on a tablet; the personal displays can be used remotely for collaboration.

ABSTRACT

One of the significant challenges of care transitions in Intensive Care Units (ICUs) is the lack of effective support tools for outgoing clinicians to find, filter, organize, and annotate information that can be effectively handed off to the incoming team. We present a large display interactive multivariate visual approach, aimed towards supporting clinicians during the transition of care. We first provide a characterization of the problem domain in terms of data and tasks, based on an observation session at the University of Illinois Hospital, and on interviews with several biomedical researchers and ICU clinicians. Informed by this experience, we design a scalable, interactive visual approach that supports both overview and detail views of ICU patient data, as well as anomaly detection, comparison, and annotation of the data. We demonstrate a large-display implementation of the visualization on an existing anonymized ICU dataset. Feedback from domain experts indicates this approach successfully meets the requirements of effective care transitions.

Index Terms: ICU care transition, large display visualization, shared cognition, collaborative decision making

1 INTRODUCTION

Intensive Care Units (ICUs) are special departments of a hospital or healthcare facility that provide critically ill or critically injured patients with intensive medical treatment, under the close supervision of specially-trained medical staff. ICUs are distinguished

from normal hospital wards by a higher staff-to-patient ratio (often a 2:1 nurse to patient ratio) [20], by access to advanced medical resources and equipment that is not routinely available elsewhere, and by access to staff with often complementary expertise. Because the ICU patients require continuous monitoring, multiple teams rotate responsibilities: each interdisciplinary team provides care for 12-18 hour shifts before they transfer the patient care responsibilities to the next critical care team. The transfer of patient-care service between an outgoing team and an incoming team is often referred to as “ICU patient handoff” or “ICU care transition” [3, 7, 14].

One of the significant challenges for care transitions in ICUs is the lack of effective tools that can help clinicians to find, filter, organize and then present patient care-related information to the incoming patient care team [4]. Two aspects of the transition process are particularly challenging: 1) gathering and integrating, in some structured format, multivariate data from multiple sources—such as vital signs from live monitoring equipment and health records that give a broader view of a patient’s medical history, resident doctor notes, or nurse notes [5] and 2) sharing the aggregated information with the incoming team such that it can be meaningfully discussed [1]. As a result, the care transition process has been critiqued as time-consuming, effortful, and prone to communication errors (35 % of medical errors) [1, 12], affecting the continuity, quality, and timeliness of patient care. Additionally, poorly executed handoffs can increase the length of stay of the patient in the ICU, incurring costs up to \$10,000 per day.

While several ICUs use spreadsheet applications, paper notes, and hard or soft copies of those notes to coordinate care transitions, advancements in health information science and technology have the potential to improve the handoff process [Sumedh S. Hoskote et al. 2016]. For example, several general mnemonics, and guidelines for the development of care transition tools have been developed in the medical field, although no standardized formats (electronic or paper-based) exist to facilitate the handoff process [Joanna Abraham et al. 2011]. At the same time, electronic displays are becoming pervasive.

*e-mail: mthoma52@uic.edu

†e-mail: tkg2@uic.edu

‡e-mail: abrahamj@uic.edu

§e-mail: gmarai@uic.edu

While the trend so far has been towards tablets and smartphones [Anthony Faiola et al. 2015] for mobility and convenience of use, doing so comes with a cognitive loss in collaborative activities like care transitions, especially in an ICU environment. With smaller displays, it is difficult to fit the information available for a patient without zooming, panning or switching the context. With respect to shared cognition, large shared displays have a significant advantage over small portable displays. Large displays facilitate teamwork by using shared contents which support data-oriented conversations. Because of these advantages, and because of the low cost of tileable displays, our ICU and biomedical informatics collaborators were particularly interested in adopting large-display technology to facilitate the care transitions process during hospital rounds.

While an increase in resolution and display size has the potential to address some scalability and information sharing challenges, novel designs are necessary to take advantage of the unique properties of these environments. Furthermore, with the huge amounts of patient data available, the potential for knowledge discovery is significant if the data is represented visually [22]. Although several visualizations of healthcare data (which includes and not limited to electronic health records, CT/MRI scans, and genomes) exist, only a handful focus on ICU data. In these ICU-related works, several multivariate visualization techniques have been proposed [6, 8, 13]. However, none of these techniques were designed to specifically support effective information organization, presentation, and collaborative decision making, and none build on the strengths of large displays.

In this paper, we describe the design and implementation of a large display interactive multivariate visual approach, which aims to support effective ICU clinician handoffs and collaborative decision making. The contributions of this work are as follows:

- A characterization of the problem domain in terms of data and tasks, based on an observation session at the University of Illinois Hospital, and on interviews with several biomedical researchers and ICU clinicians.
- The design, in collaboration with domain experts, of a multivariate interactive visual approach that supports both overview and detail views of ICU patient data. The visualization supports anomaly detection, comparison, and annotation of the data.
- A large display implementation and demonstration of the resulting application (Echo: Effective Care HandOffs) on an existing anonymized ICU dataset, with feedback from domain experts, and a discussion of the experts findings and observations.

2 RELATED WORK

Electronic Health Record (EHR) Visualization. EHR-based patient records tend to be large, complex, and disorganized which makes its exploration and insight-finding difficult, impacting patient care. To overcome such limitations, LifeLines [16] support details on demand, a zoomable timeline, relationship highlights, outlining and summarizing. Even though our approach follows similar principles of information mapping, we avoid the zoom and panning approach to preserve contextual information. LifeLines2 [21] uses an alignment technique to compare similar patients data and predict a possible event. In contrast, our focus is on summarizing multiple patient records to select patients based on specific criteria (either newest patient or critical patient). VisuExplore [17] uses line charts and bar charts to show both numerical and categorical data on one screen. These simple visual encodings helped clinicians at multiple ranks (residents, fellows, interns). The design of our tool also consists of line charts to represent numerical data since the tool is intended for an ICU team comprised of medical staffs

at different ranks, including nurses. A web-based approach called WBIVS [15] uses similar simple encodings for categorical variables, as well as average lines over multiple time steps, to reduce visual clutter. TimeRider [18] uses animation to represent a chance to show trends in a cohort. Animations can be distracting when the number of variables is high, even though animation traces are useful in identifying data overlaps. Our large display approach overcomes the contextual loss caused by animations and avoids the visual clutter caused by traces or multi-step lines. It is also worth noting that most of these systems were designed for secondary uses of data (i.e., retrospective), as opposed to direct, EHR-based, prospective use during hospital rounds.

Temporal ICU data visualization. Many techniques focus on showing temporal ICU data through stack plots of univariate time series graphs. One such technique is the Medical Information Visualization Assistant (MIVA) [6], which uses stacked point plots to visualize critical care patient data; records can be added or reordered using a drag and drop interaction. The stacked graphs are harder to read as the number of graphs (variables) increases, and the user has to scan through the interface to see the trends. Additionally, MIVA shows the records of only one patient, and there is no support for comparing patients, or for viewing summaries of all the patients to group them. The MTSA (Multivariate Time Series Amalgam) system [13] takes a different approach to overcome clutter, which makes use of radar charts and animations to display vital signs and their change over time. The use of animation can be distracting, and due to the amount of data, the interface can get crowded and difficult to interpret. Another technique for single patients called metaphor graphics [9] shows vital signs as small glyph multiples, where each multiple corresponds to one hour. The small multiples are useful in maintaining contextual information, but the glyph that changes size, shape, and color can be confusing. A recent technique, Care Cruiser [8] visualizes the effects of treatment plans, in addition to the patient data, by using color-coding to show noteworthy patient events. However, Care Cruiser and other display techniques surveyed in this section were not designed for and do not support team communication during care-transition meetings.

Display Technology and EHR. Advances in display technology have recently made an impact on EHR visualization. An et al [2] follow a dynamic approach inspired by the level of detail (LOD) rendering used in 3D computer graphics. The context switching can make it harder to see how different data are connected. New display technologies and handheld devices are replacing the old CRT monitors in the ICU, with tablets, in particular, becoming popular in the medical field due to their portability and ease of use. A Watson-aided visualization of EHR data [19] uses machine learning to select only the relevant information, which can be navigated further down if necessary. However, personal displays do not establish a common information space to help a health-care team share insights, to develop an optimal treatment plan.

3 METHODS

3.1 Domain Characterization

We gathered information related to the ICU health care transition process by conducting a series of interviews with two medical researchers and a clinician. The interviews were followed by an observation session at a medical ICU in an academic medical center, involving an incoming and an outgoing team that comprised attending physicians, fellows, residents, interns and nurses, and optional specialists and pharmacists. Unlike general hospital multidisciplinary rounds, on which there is ample research available, the ICU care transitions process is a less well understood, one-time event.

Diagnosis	Physical Exam/Labs	Medications	Problem list	Resident Plan	Attending Plan
NEUROLOGIC <input type="checkbox"/> Delirium <input type="checkbox"/> Coma <input type="checkbox"/> Encephalopathy <input type="checkbox"/> Anoxic Brain Injury <input type="checkbox"/> Critical Illness Polyneuropathy <input type="checkbox"/> Cerebrovascular Accident/Stroke <input type="checkbox"/> Seizure <input type="checkbox"/> Syncope <input type="checkbox"/> Other	Psych: <input type="checkbox"/> Sedated/ <input type="checkbox"/> Agitated/ <input type="checkbox"/> Calm Neuro: <input type="checkbox"/> Sedated/ <input type="checkbox"/> Confused/ <input type="checkbox"/> Alert-Awake-Oriented_ <input type="checkbox"/> Cinnial Nerves intact, Motor strength in <input type="checkbox"/> <input type="checkbox"/> UE <input type="checkbox"/> LE 5/5 and symmetrical. <input type="checkbox"/> <input type="checkbox"/> DTRs 2+ and symmetrical	Sedatives: <input type="checkbox"/> Lorazepam <input type="checkbox"/> Midazolam <input type="checkbox"/> Fentanyl <input type="checkbox"/> Propofol <input type="checkbox"/> Precedex	<input type="checkbox"/> Awake Alert <input type="checkbox"/> Altered Mental Status	<input type="checkbox"/> Daily Interruption of Sedation	
CARDIOVASCULAR <input type="checkbox"/> Shock, unspecified <input type="checkbox"/> Septic Shock <input type="checkbox"/> Cardiogenic Shock <input type="checkbox"/> Hypotension <input type="checkbox"/> Hypertension, stage I <input type="checkbox"/> Hypertension, stage II <input type="checkbox"/> Myocardial Infarction <input type="checkbox"/> Systolic Heart Failure <input type="checkbox"/> Diastolic Heart Failure <input type="checkbox"/> Valvular Heart Disease <input type="checkbox"/> Coronary Artery Disease <input type="checkbox"/> Atrial Fibrillation <input type="checkbox"/> Supraventricular Tachycardia <input type="checkbox"/> Sinus Tachycardia <input type="checkbox"/> Cardiac Arrest	BP / MAP HR CVP CV: Rhythm: <input type="checkbox"/> Regular/ <input type="checkbox"/> Irregular, Rate: <input type="checkbox"/> Bradycardic/ <input type="checkbox"/> Normal/ <input type="checkbox"/> Tachycardic/ <input type="checkbox"/> Rub Murmurs: <input type="checkbox"/> Yes <input type="checkbox"/> No; <input type="checkbox"/> Systolic/ <input type="checkbox"/> Diastolic; Grade: <input type="checkbox"/> VI, <input type="checkbox"/> Location: _____; <input type="checkbox"/> Radiation _____	Vasopressors: <input type="checkbox"/> Norepinephrine <input type="checkbox"/> Dopamine <input type="checkbox"/> Dobutamine <input type="checkbox"/> Phenylephrine <input type="checkbox"/> Vasopressin Antihypertensives <input type="checkbox"/> Beta-blockers <input type="checkbox"/> Calcium channel blockers <input type="checkbox"/> ACE-inhibitors			

Figure 2: Progress note using a SOAP (Subjective, Objective, Assessment, Plan) template prepared by an outgoing clinician and used by the team for discussion.

3.1.1 Interview Analysis

We conducted multiple interviews with three domain experts: two biomedical informatics researchers (with expertise in ICU Communication and Health Information Technologies) and one practicing ICU clinician.

The interviews indicated that two critical care teams (on-call or incoming, and post-call or outgoing) participate in the care transition process; there is typically no overlap between the two teams. During shift changes, the teams meet face-to-face for the handoff. The discussion per session consists of sixteen to eighteen patients, with as little as 5-7 minutes allocated per patient, for a total duration of 2-3 hours.

Before the meeting, the outgoing resident physician responsible for the patient collects data regarding the patient and enters it into a critical care progress note form. While the format of this progress note may vary across hospitals, its content is often standardized. The progress note used at this facility is based on a problem-based SOAP (subjective objective assessment and plan) format as shown in Figure 2. For each body system, the form lists a Diagnosis, Exams, Medications, and a Problem list. An additional “plan section” records the proposed plan of care. The Diagnosis section features a list of predefined diagnosis for each body systems, on which the outgoing clinician marks the ones which are relevant to the patient. Physical/lab exam results and the waveform data which continuously monitors the patient are crucial in making the diagnosis. The Physical/Lab Exam section features a predefined exam field, which the clinician fills in with the data for the exams performed. The Medications section lists all fluid or oral drugs that were given to the patient for each body system in the past 24 hours. Finally, the Problem section lists all the problem associated with each body systems due to medication or any other factors.

During the care-transition, the completed progress note drives the presentation and discussion. The outgoing resident presents his/her notes, with the rest of the team contributing additional information or clarification questions. Along with the transfer of information, authority and responsibility for the patient, a patient plan of care for the next 24h are also developed as part of this process. Additional outcomes of the discussion are an assessment of whether the patient status has improved or not, and the level of team agreement with the plan of care and patient status.

The domain experts were in agreement that the most challenging aspect of the care transition process was sharing the aggregated information to the incoming team which supports meaningful discussion.

3.1.2 Observation Session

Following the interviews, we performed an observation session in the MICU (Medical Intensive Care Unit) at the University of

Illinois Hospital, to verify the interview findings, to get a better understanding of how clinicians and their teams interact with the data, and to note how they communicate and collaborate in the clinical decision process about the treatment of a patient. The MICU is an 8-bed unit that provides care for adults and geriatric patients, in which single-bed rooms line down a wall. The environment was noisy, with many alarm systems going on and off frequently; a dispatch and monitoring center was located across the rooms. These observations endorsed the clinicians request for a salient visual communication channel. A set of flat panels was mounted on the wall above the room entrances, further supporting the experts’ interest in tiled displays. However, the panels only displayed basic information about each bed, such as the patient name and basic vitals.

In our observation session, a large team of twelve people (including the outgoing and incoming teams) started the round session from one end of the ICU and proceeded to the other end, each bed taking around 20 minutes. The resident in charge pointed out that some teams prefer discussing in a conference room, while this team preferred to be able to visit the patient during the discussion if need be; however, she indicated that the structure of discussion is the same regardless of the handoff setting. The discussion started with the outgoing clinician handing out to both team members copies of the progress note she had prepared for each patient. For each patient, the outgoing resident then gave a brief presentation based on the prepared progress note, followed by questions and discussion with the incoming team.

If there was a question involving data which was not available in the progress note, the team extracted the required data from a computer system. In this case, the system was a computer mounted on a stand with wheels called CoW (computer on wheels), that the team pushed around as they progressed down the hallway. The CoW typically contains all the available data of a patient, usually in a spreadsheet-like format; however, the CoWs relatively small display was only accessible to a couple of team members at a time. Interaction among team members was reduced, with most participants studying the notes in their hands, and the clinician in charge reading out those notes. Nurses, who routinely monitor the patients, provided occasional additional information based on their written notes. A dispatch call was occasionally made to the pharmacist to clarify a potential medication issue. The clinician in charge entered briefly the room to further examine the patient only in the case of two patients, while most of the team waited outside; otherwise, all discussions took place in the hallway, based on the patient notes and the CoW. Based on each discussion, the team prepared a treatment plan for the patient. The plan was then assigned to the incoming care team to carry out.

While the discussion process was amenable to technological improvements, the teams we observed reported during the debriefing that the most challenging aspect of the process was gathering the data and preparing it for presentation to the team. This finding was in some contradiction with the earlier interview findings, and lead us to design an approach that supports both gathering the data and the discussion process.

In terms of requirements, we further learned that:

- To create a summary of all the patient in the ICU, the outgoing clinician has to go through a time-consuming process of manually scanning through all the available data, which resides in a database of an EHR system. The outgoing clinician has to extract the actionable items from a significantly large amount of data, and the chance of missing vital information is high.
- The teams wished for a way to determine the appropriate order in which to discuss the patients: newest patients first, or most critical patients first. According to their experience, the patient case discussed later will get less time and attention.

- In the discussion stage, the teams rely on the CoW to get any missing or previous data. However, the CoW is only capable of providing raw values, without any further insights. No visual encodings are used to support the discussion, which makes it hard to detect trends, outliers, or correlations between, for example, today's data and yesterday's data.
- Beyond the progress note, there is no pool of shared information among the team, since each member concentrates on their copy of the notes. This creates opportunities for cognition loss and errors.

3.1.3 Data Types

The data accessed during the care transitions process includes vital signs over the past 24 hours such as heart rate, pulse, temperature, respiratory rate, mean arterial pressure (MAP) and oxygen saturation rate (SpO₂). These data are recorded in numerical form every minute and are typically stored in the waveform. Other data considered are chart events in the form of physical/lab exam results or medication details. These data are discrete, and consist of both numerical and text fields, and they may be grouped into specific categories. The data may also include summaries such as a morbidity score, which is used to rate patients based on the severity of the condition, the medical condition, and the current status that shows whether the patient is improving, stable or critical.

Due to HIPAA regulations, our current implementation uses data from a publicly available dataset called MIMIC-III. This large dataset contains multi-dimensional de-identified healthcare data associated with over 60,000 patients who stayed in critical care units of the Beth Israel Deaconess Medical Center between 2001 and 2012. MIMIC-III provides high temporal resolution of data, including lab results, electronic documentation, and bedside monitor trends and waveforms. We extracted three days worth of data of several patients, consisting of waveform numerics of two body systems, cardiovascular and pulmonary, as well as chart events including medications. Because the waveform information is recorded for every minute, while the experts specifically stated they were interested in hourly values (brief spikes in values are considered insignificant), we computed an hourly aggregate of these waveform data.

3.2 Visual Design

Our domain characterization shows that both 1) gathering actionable data and 2) presenting the information plays a major role in the care transition process. Additionally, 3) the ability to group and prioritize patients can help a team manage its efforts more efficiently. Based on this analysis, our design supports these three essential aspects of the handoff process. Our final top-level design comprises two displays: a summary of all patients in the ICU (aspect 3), and a detail view for an individual patient (aspects 1 and 2). This visual approach was developed through an Activity-Centered-Design approach (an extension of Human-Centered-Design that emphasizes user activities and functional specs) with parallel prototyping, in which frequent meetings with and repeated feedback from two of the domain experts drove the design decisions. We describe below each of the two resulting displays.

Because we aim to support collaborative decision making with a high degree of awareness among groups, we developed our tool using the SAGE2 [11] application API. SAGE2 is a windowed operating environment for large-scale shared displays that provides multiple forms of remote collaboration, and simultaneous multi-user interaction using mouse pointer and, where available, direct touch. On the hardware side, SAGE2 works on any large display or low-cost tileable displays. On the software side, SAGE2 is web-based, which allows the building of apps or visualizations using any web technologies currently available on the market. The resulting applications can be ported to work on any normal browser with minimum

changes. In this case, we used d3.js for all the visualizations and snapsvg for all the vector graphics.

3.2.1 Summary overview of all the patients in the ICU

Because our domain characterization showed that the discussion order of the ICU cases is particularly important during care transitions, the entry-point to our visualization is a daily summary of all the patients in the ICU. The summary can also be used by clinicians and nurses to monitor the status of a patient routinely. Based on this overview, patients can be easily identified according to a teams priorities, for instance discussing first those patients who are in a critical condition.

To take advantage of the display space, we show the patient summaries in a small multiple card layout as shown in Figure 3. Each card contains identifying information about that patient (in MIMIC-III, de-identified name, age, and a placeholder picture), a composite patient status, a small set of icons indicating medical conditions, and a summary showing the deviation of vital signs from the normal range. The card layout takes advantage of the large display by placing all the cards on the screen at the same time.

Summary Statistics: According to the domain experts, the length of stay in the ICU, the number of days with assisted ventilation, along with changes in vital signs, the main reason for the ICU admission, the morbidity score and the state of a patient gives the most important summary information about a patient. To compute these measurements, we took the aggregate of all the vital signs, separately, for the past 24 hours, and normalized those values into a 1-10 range, with one being the lowest. This approach enabled us to calculate how much these values deviate from normal, and thus prioritize them. In addition to the vital signs, for each patient, we show the medical condition (such as stroke, lung failure, heart failure, renal failure, liver failure or trauma) which was responsible for the ICU stay.

The summary statistics are a key component of the interface and are encoded using a Kiviat diagram as shown in Figure 5. In the Kiviat, the normalized vital sign information is mapped to the Kiviat axes, showing how much each vital parameter deviates from the normal range. Dots indicate normal ranges on each axis. The Kiviat color is mapped to the morbidity score: low saturation corresponds to less severe conditions, and high saturation corresponds to higher severity and an increased risk of death. We selected the Kiviats over other encodings (such as parallel coordinate plots) that can be confusing to clinicians and nurses, whose visual literacy is typically low. The Kiviats have the advantage of making quick comparisons possible.

Composite Patient Status: A triangular glyph on each card indicates whether the patient is improving, stable or critical. To encode the status, the vital signs and the morbidity score of the current day are compared with the previous day. The glyph takes a gray color if all the values are in the normal range, blue if the values show improvement, respectively orange if values deteriorate. The outgoing clinician can use this indicator to pick the patients who are critical and discuss their case with the team. Body system icons are placed on each card to show the failing organs or diseases which caused the admission of the patient to the ICU.

Clinicians can navigate to each patient's detailed view by clicking on that patient card in the summary display. If the display has touch support, touching the card will bring up the detailed view.

3.2.2 Detailed view of an individual patient

The detailed view of a single patient displays waveform vitals and exam information, chart events, and medication details, along with the patient identification at the top as shown in Figure 4. Because of documented low visual literacy among nurses [10] and similar literacy among the rest of the team, this display includes primarily basic visual encodings such as plots, charts, and recognizable icons,



Figure 3: Card layout, showing the summary overview of 8 patients in the ICU. Each card consists of a Kiviat plot showing summary statistics of vital signs, a morbidity score, the medical condition, and a composite patient status. The displays show at a glance that the most critical patients in this set are John Doe, Anna Ross, and Joan Lee (dark purple Kiviat shades). The status of all these three patients is also worsening (orange status glyph). John Doe (top left) has been longest on respiratory support (Vent: 1) and should be discussed first. Underneath, Ann Cook has been longest both in the ICU and on respiratory support, but her status is stable.

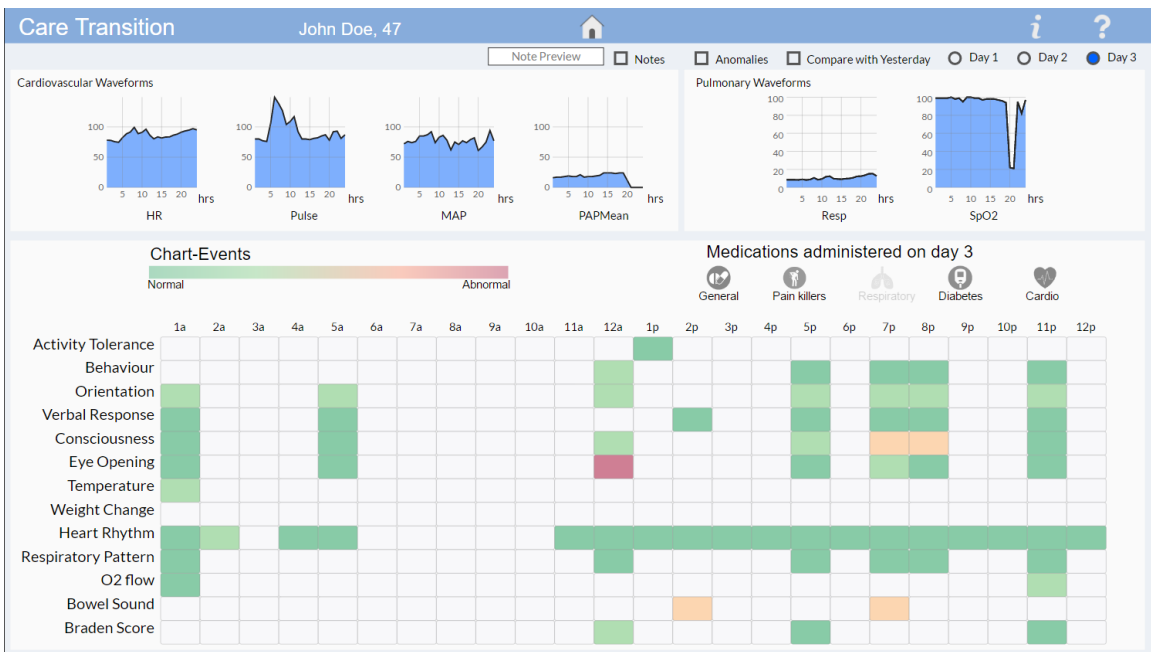


Figure 4: Detailed view of an individual patient, showing 1) waveform data for two body systems: cardiovascular and pulmonary (top), 2) chart-event heatmap of 13 most frequently checked events (bottom), and 3) medications given on that day (middle), along with 4) options to show notes, anomalies, comparison and previous days data.




Figure 5: (Left) Kiviatic plot showing summary statistics, with normal values indicated by black dots. The Kiviatic dark shade indicates high morbidity score. (Right) Kiviatic legend, with the axes mapped to: MAP (mean arterial pressure), Resp (respiratory rate), HR (Heart Rate), SpO2 (Blood oxygen saturation level), and Temp (temperature)

augmented by text and numerics. To compensate for the reduced expressive power of these encodings, as well as to support the content gathering and presentation of patient data, we augment this display with capabilities for automated anomaly detection, for data comparison, and for annotation. The overall design of this display follows the structure of a patient progress note.

For demonstration purposes, the display includes the only two body systems in the MIMIC database - Cardiovascular and Pulmonary. The parameters we show for the Cardiovascular system are heart rate, pulse, mean arterial pressure and mean pulmonary artery pressure, and for Pulmonary respiratory rate and oxygen saturation rate. Because chart events are discrete data consisting of both numerical and textual values, we quantify each value with a “normalcy score index”. We show the following 13 most frequently checked chart events, sorted by body system: activity tolerance, behavior, orientation, verbal response, consciousness, eye-opening, temperature, weight change, heart rhythm, respiratory pattern, O2 flow, bowel sound, Braden score. Medications are grouped into five categories: general, pain-killers, respiratory, diabetes and cardiovascular.

Vitals: We show the hourly average of each parameter of both body systems as small multiples of line charts. Line charts are suited for encoding waveform data as they show temporal variation, which in turn helps to identify trends. Using a small multiple display not only avoids overplotting of parameters in a single chart but also allows us to compare different parameters and highlight potential anomalies. We further use overlays of the line charts to enable comparison with previous days data without losing context.

Chart Events: The chart-events are represented as a pixel heat map as shown in Figure 6, where each row represents different events, and each column the time (24 hours). The augmented normalcy score index is used to show whether the event had a normal or an abnormal outcome. A hover action can be used to show on demand the logged value for each chart-event. The pixel heat map provides a complete and compact picture of all the events associated with a patient and illustrates the connections between the events. The pixel heat map was selected over other encodings because, according to our domain experts, hospital personnel are used to grid-like visual encodings, which have a tabular chart structure.

Medication: Because medication entries in the MIMIC database are associated with days, not with hours, we use standard glyphs to encode the type of medication given to the patient. The glyphs associated with medication classes not applied are grayed out. Hovering over the medication group will show the list of medications given on that day in an expanded view.

Users can navigate in this display through up to 3 days of data. A radio button for each day allows the user to navigate through consecutive days and see the waveforms, chart-events, anomalies, and annotations associated with each day. Three toggleable checkboxes show/hide anomalies, comparison, and notes.

Anomaly Detection: Because the normal and abnormal values for parameters of each body system vary from patient to patient, they get assigned in the EHR by the attending physician. To detect anomalies, we compare the average value of each parameter against the assigned normal. We highlight the corresponding point in the line chart as shown in Figure 7 (Left).

Supporting Comparison: Waveform data of two consecutive days can be compared by overlaying the line charts as shown in Figure 8. This approach allows the clinician to see how each parameter is responding to the treatment plan set for that particular day when compared to the previous treatment plan.

Supporting Annotations: The annotation capability allows clinicians and nurses to add observations that were made during routine monitoring or exam of the patient. The annotations can be placed through a point and click action anywhere on the display, for example next to the event or waveform referred by the note as shown in Figure 7 (Right). A note preview shows the note content before publishing. The annotations can be expanded on demand, allowing the outgoing clinician to use this information during the presentation to the team.

4 EVALUATION AND RESULTS

We tested multiple versions of our approach with three naive users, as well as with two of our collaborators, who are active researchers in the fields of health informatics and family medicine. The evaluations were done on a variety of displays, from laptop displays to a 74.5-inch screen, and to a wall-sized display environment.

Naive Testers: The naive testers were given an objective and a scenario in which each tester was to take the role of the outgoing clinician. The tester was tasked with gathering relevant information to be presented before a fictional team. Four tasks were given, along with a basic description of all the parameters used in the interface:

- Task 1 - Select the patient whose case should be discussed first.
- Task 2 - Find at least two parameters which need to be discussed with the team.
- Task 3 - Compare waveforms with the previous day data and see if there is a trend.
- Task 4 - Find out whether there is a correlation between the cardiovascular and pulmonary system data

The testers were given full control over the interface. Each tester took 5-7 minutes to complete the tasks given. Even though they lacked medical experience, they made similarly sound assumptions and choices that were close to the ground truth. The testers were able to detect outliers, trends, and correlations in the data. An open-ended questionnaire followed the testing session. An earlier version of the tool had raw values mapped to the Kiviatic axes, which caused confusion among the testers—for some parameters normal was indicated by a high value, and for others by a low value. To tackle this issue, we normalized the values so that one is always the lowest. To compare waveforms between two days, users had to switch back and forth between days and had difficulty remembering features. This motivated us to implement an overlay design, which allows comparison without switching context.

Expert Feedback: The two domain experts (medical informatics experts with critical care research experience) appreciated the novelty of using the large display in the ICU data intensive environment and its usefulness to communication between teams. The potential use of the approach to both gathering of data and presentation to the team was immediately noted. They further found the patient card layout to be useful not only for routine monitoring but also when organizing the patients based on 1) morbidity score, 2)




Figure 6: Chart-events showing 13 frequently checked chart-events, and hover interaction showing the value that was logged at that time. The chart-map raises multiple red flags with respect to the patient status throughout the day.




Figure 7: (Left) Automatic detection of pulse waveform anomalies highlighted in red. The average value per hour is compared with the normal value assigned to the patient; if the value is outside the normal range, that portion is highlighted with a red marker. (Right) Toggleable annotation placed next to the HR (heart rate) waveform.




Figure 8: (Left) Comparison with previous days data for a Pulse waveform. The blue graph indicates the current day data, and the gray graph shows the previous day data. (Right) Pulse waveform with both the comparison and anomaly detection options turned on. The current day anomalies are highlighted in red, and the previous day is in yellow.

organ failures and 3) stability. They specifically commented that the card approach would help to bring in optional specialists efficiently. The Kiviat encoding was easily deciphered: the experts were immediately able to identify which cases were in worse condition. They further found the triangle glyph on the cards to be a good indicator of whether the current treatment plan was effective or not.

The collaborators liked the organization and grouping of data in the detailed view display. They commented that there was no existing system to group and visualizes this amount of data under a single view, which was important in maintaining context. They found the waveform representation and the options associated with it, and the chart event map to be useful in the detection of trends and anomalies in the data. They liked the annotation capability and the fact that the annotations were toggleable, and that they could be placed anywhere on display. The comparison capabilities allowed them to further compare multi-day data without switching views. They also liked the fact that all options were toggleable, which avoided information clutter. The experts felt this further made the system a good potential candidate for training purposes.

5 DISCUSSION AND CONCLUSION

The evaluation with naive users and with domain experts shows our tool to be simple enough and powerful for handoffs and collaborative discussions. The scalable display visualization has further shown promise in terms of showing information without cluttering or loss of context. This large-display visualization helps create a salient visual communication channel, which can be easily accessed by all team members, whether they are co-located or remotely. According to the expert feedback, these aspects can help in creating shared

cognition among a care transitions team. Through a participatory design approach, we confirmed the importance of simple, familiar visual encodings to healthcare practitioners.

As shown by our evaluation, the two displays we have created successfully meet the original requirements of the target domain. The first display shows daily summaries all the patients in an ICU, which helps clinicians to organize and select the most urgent cases first. The detailed view shows all the body system information of an individual patient. This second view supports, through anomaly detection, comparison, and annotations, the gathering, and presentation of relevant information. Even though the visualization was designed to support efficient handoffs and decision making, it is also capable of supporting routine monitoring of patients. The expert feedback further identified the potential of this work to training junior clinician, nurses or interns, in helping these trainees find actionable items and communicate them to a team.

In terms of assumptions and limitations, due to HIPAA regulations, our implementation uses a de-identified dataset with a limited number of body systems and data fields. Because our approach was designed from the start for a large display environment, we nevertheless expect the visualization to scale reasonably well with an increasing number of measurements. Our work further assumes the EHR data is already stored in electronic format, while in reality some of the measurements might be streaming in real time from medical devices. We made this assumption based on our domain process analysis, in which residents were gathering and presenting pre-collected data.

In conclusion, we have designed and implemented a large display

multivariate interactive visualization of ICU data for healthcare handoffs and collaborative decision making. Our visual approach builds on our characterization of the problem domain, which we performed with experts in the target domain, through interviews and direct observation of their workflow. The two-view design of our approach successfully meets the target requirements for grouping patients, gathering data and presenting information in a salient visual channel, and further supports anomaly detection, comparison, and annotation of data. The evaluation with naive testers and feedback from domain experts shows this approach can be useful in facilitating more efficient care transitions in the ICUs.

ACKNOWLEDGMENTS

This work has been partially supported by UIC College of Engineering Award “Collaborative Dashboard for Patient Care Transitions”, and by National Science Foundation Awards NSF IIS-1541277 and NSF CNS-1625941. We gratefully acknowledge the support of the Electronic Visualization Lab, in particular J. Trelles, A. Burks, and L. Long, and the generous assistance of Prof. Min J. Joo, University of Illinois Hospital.

REFERENCES

- [1] J. Abraham, V. Nguyen, K. F. Almoosa, B. Patel, and V. L. Patel. Falling through the cracks: information breakdowns in critical care handoff communication. In *AMIA Annu Symp Proc*, vol. 2011, pp. 28–37, 2011.
- [2] J. An, Z. Wu, H. Chen, X. Lu, and H. Duan. Level of detail navigation and visualization of electronic health records. In *Biomedical Engineering and Informatics (BMEI), 2010 3rd International Conference on*, vol. 6, pp. 2516–2519. IEEE, 2010.
- [3] V. Arora, J. Johnson, D. Meltzer, and H. Humphrey. A theoretical framework and competency-based approach to improving handoffs. *Quality and Safety in Health Care*, 17(1):11–14, 2008.
- [4] L. Cheah, D. H. Amott, J. Pollard, and D. A. Watters. Electronic medical handover: towards safer medical care. *Medical Journal of Australia*, 183(7):369, 2005.
- [5] P. R. Ebright, L. Urden, E. Patterson, and B. Chalko. Themes surrounding novice nurse near-miss and adverse-event situations. *Journal of Nursing Administration*, 34(11):531–538, 2004.
- [6] A. Faiola, P. Srinivas, and J. D. Duke. Supporting clinical cognition: A human-centered approach to a novel icu information visualization dashboard. *AMIA ... Annual Symposium proceedings. AMIA Symposium, 2015:560–9*, 2015.
- [7] D. Goldsmith, M. Boomhower, D. R. Lancaster, M. Antonelli, M. A. M. Kenyon, A. Benoit, F. Chang, and P. C. Dykes. Development of a nursing handoff tool: A web-based application to enhance patient safety. In *AMIA Annu Symp Proc*, vol. 25, pp. 6–60, 2010.
- [8] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. Design and evaluation of an interactive visualization of therapy plans and patient data. In *Proceedings of the 25th BCS Conference on Human-Computer Interaction*, pp. 421–428. British Computer Society, 2011.
- [9] W. Horn, C. Popow, and L. Unterasinger. Metaphor graphics to visualize icu data over time. *Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-98)*, 1998.
- [10] K. D. Lopez, D. J. Wilkie, Y. Yao, V. Sousa, A. Febretti, J. Stifter, A. Johnson, and G. M. Keenan. Nurses’ numeracy and graphical literacy: Informing studies of clinical decision support interfaces. *Journal of nursing care quality*, 31(2):124–130, 2016.
- [11] T. Marrinan, J. Aurisano, A. Nishimoto, K. Bharadwaj, V. Mateevitsi, L. Renambot, L. Long, A. Johnson, and J. Leigh. Sage2: A new approach for data intensive collaboration using scalable resolution shared displays. In *Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on*, pp. 177–186. IEEE, 2014.
- [12] K. P. Mistry, J. Jaggars, A. J. Lodge, M. Alton, J. M. Mericle, K. S. Frush, and J. N. Meliones. Using six sigma® methodology to improve handoff communication in high-risk patients. 2008.
- [13] P. Ordóñez, M. desJardins, M. Lombardi, C. U. Lehmann, and J. Fackler. An animated multivariate visualization for physiological and clinical data in the icu. In *Proceedings of the 1st ACM International Health Informatics Symposium*, pp. 771–779. ACM, 2010.
- [14] E. S. Patterson, E. M. Roth, D. D. Woods, R. Chow, J. O. Gomes, et al. Handoff strategies in settings with high consequences for failure: lessons for health care operations. *International Journal for Quality in Health Care*, 16(2):125–132, 2004.
- [15] D. S. Pieczkiewicz, S. M. Finkelstein, and M. I. Hertz. Design and evaluation of a web-based interactive visualization system for lung transplant home monitoring data. In *AMIA Annual Symposium Proceedings*, vol. 2007, p. 598. American Medical Informatics Association, 2007.
- [16] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, and B. Shneiderman. Lifelines: using visualization to enhance navigation and analysis of patient records. In *Proceedings of the AMIA Symposium*, p. 76. American Medical Informatics Association, 1998.
- [17] M. Pohl, S. Wiltner, A. Rind, W. Aigner, S. Miksch, T. Turic, and F. Drexler. Patient development at a glance: An evaluation of a medical data visualization. *Human-Computer Interaction—INTERACT 2011*, pp. 292–299, 2011.
- [18] A. Rind, W. Aigner, S. Miksch, S. Wiltner, M. Pohl, F. Drexler, B. Neubauer, and N. Suchy. Visually exploring multivariate trends in patient cohorts using animated scatter plots. In *International Conference on Ergonomics and Health Aspects of Work with Computers*, pp. 139–148. Springer, 2011.
- [19] P. Ruchikachorn, J. J. Liang, and K. Mueller. Watson-aided non-linear problem-oriented clinical visit preparation on tablet computer. 2014.
- [20] J. Wang, V. Lakticova, A. M. McCabe, B. Amitrano, L. Rosen, and R. I. Cohen. Critical care nurses perception of time spent at rapid responses. *Annals of the American Thoracic Society*, 10(3):228–234, 2013.
- [21] 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 *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 457–466. ACM, 2008.
- [22] V. L. West, D. Borland, and W. E. Hammond. Innovative information visualization of electronic health record data: a systematic review. *Journal of the American Medical Informatics Association*, 22(2):330–339, 2015.