

Sanguine: Visual Analysis for Patient Blood Management

Haihan Lin*
University of Utah

Ryan A. Metcalf†
ARUP Laboratories, University of Utah

Jack Wilburn‡
University of Utah

Alexander Lex§
University of Utah

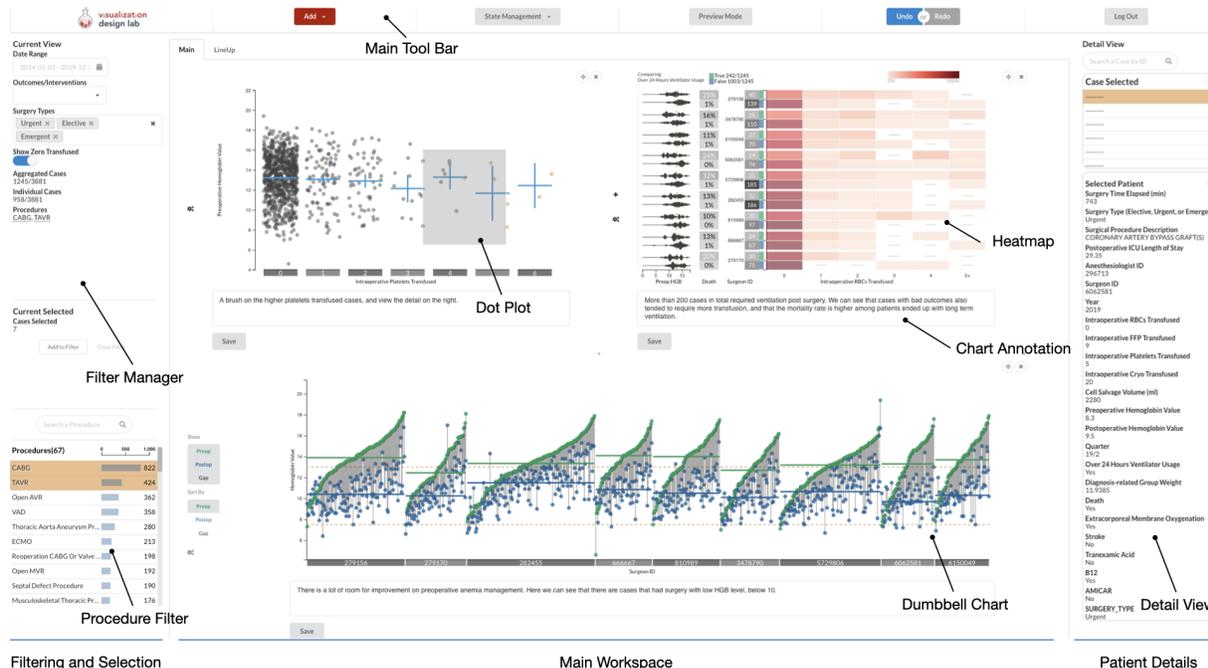


Figure 1: An overview of Sanguine visualizing patient blood management data with multiple views. The left panel is dedicated for managing filters and selections. A center workspace contains visualizations that can be flexibly arranged. A heatmap, a dot plot, and a dumbbell chart are shown. On the right, a patient-specific detail view shows attributes of a selected case.

ABSTRACT

Blood transfusion is a frequently performed medical procedure in surgical and nonsurgical contexts. Although it is frequently necessary or even life-saving, it has been identified as one of the most overused procedures in hospitals. Unnecessary transfusions not only waste resources but can also be detrimental to patient outcomes. Patient blood management (PBM) is the clinical practice of optimizing transfusions and associated outcomes. In this paper, we introduce Sanguine, a visual analysis tool for transfusion data and related patient medical records. Sanguine was designed with two user groups in mind: PBM experts and clinicians who conduct transfusions. PBM experts use Sanguine to explore and analyze transfusion practices and its associated medical outcomes. They can compare individual surgeons, or compare outcomes or time periods, such as before and after an intervention regarding transfusion practices. PBM experts then curate and annotate views for communication with clinicians, with the goal of improving their transfusion practices. Such a review session could be in person or through a shared link. We validate the utility and effectiveness of Sanguine through case studies.

Index Terms: Human-centered computing—Visualization—Visualization application domains—Information visualization

*e-mail: hmlin@sci.utah.edu

†e-mail: ryan.metcalf@path.utah.edu

‡e-mail: jwilburn@sci.utah.edu

§e-mail: alex@sci.utah.edu

1 INTRODUCTION

Transfusion is the most commonly performed medical procedure during hospitalizations in the United States [25], with 17.2 million blood products transfused in 2015 [6]. Although transfusions are often medically necessary and even life-saving, they have also been identified as one of the most overused procedures in hospitals [33]. Unnecessary transfusions can lead to adverse reactions and complications, such as acidosis, hypothermia, coagulopathy, and potentially even death [30].

Management and acquisition of blood products is also expensive, creating financial incentives for hospitals to optimize their usage. To better address these problems, many hospitals have hired dedicated experts in patient blood management (PBM), and PBM has become an independent academic field. Through PBM, health care providers aim to reduce unnecessary transfusions and improve patients' outcomes through various methods, such as treating anemia (lack of red blood cells) prior to surgeries, providing guidelines on when transfusions are necessary, and using cell salvage ("recycling" blood) during surgery [18].

Successful PBM involves multiple stakeholders. PBM experts analyze the usage of blood products and give advice or develop guidelines on best practices. Clinicians, in turn, make decisions about when and how much to transfuse, and hence they need to be well informed about best practices and their individual performance regarding these best practices.

These analysis processes, however, are not possible with the current tools available in hospitals. Electronic health records software may provide static charts and reports on blood product usage, but they lack the nuance and sophistication necessary to make a holistic evaluation of the practice of individual clinicians. For example, when one surgeon uses significantly more blood products than their

peers, they might rightly claim that they do so because they are treating the most difficult cases with many comorbidities.

To tackle this issue, we created Sanguine, an interactive visual analytic tool for PBM analysis. Sanguine was developed as a design study using an iterative design process involving PBM experts, surgeons, and anesthesiologists at the University of Utah Hospital. Sanguine enables analysts to view outcomes juxtaposed with blood product usage data, review practices on patient blood management of different providers, and compare the effect of patient blood management in clinical settings. Sanguine is designed with two types of stakeholders in mind: PBM experts who conduct in-depth analysis, and clinicians, who receive tailored and annotated interactive reports from the PBM experts. We designed Sanguine to facilitate communication around patient blood management.

Our main contribution is the design and development of Sanguine, an open-source visual analytic tool for the analysis of patient blood management data. We also contribute a domain task analysis and task abstraction. We evaluate the utility of Sanguine in case studies and through interviews with domain experts.

2 MEDICAL BACKGROUND

Patient blood management encourages minimizing blood use through the appropriate use of blood products, ascertained by evidence-based transfusion strategies [15, 31]. However, the decision to transfuse often depends on the clinical experience of the practitioners, and blood product usage can vary significantly among individuals. Analysis of transfusion records [9, 10] has shown how different surgeons within the same department vary in their blood use. Current electronic medical record systems do not give practitioners the opportunity to easily review and reflect on multiple cases simultaneously, and the static charts and reports used by clinicians to review transfusions limit them from exploring the collected data and reflecting on detailed cases.

Blood banks at hospitals regularly report wasted blood products, primarily because of excessive ordering of blood products during preoperative preparations. A blood product issued to the operating room, if unused, must be returned to the blood bank in a timely manner and properly stored; otherwise, it must be discarded. Estimates of wasted blood components range from 200,000 to 1 million in the US annually, with associated costs of 46 million to 230 million US dollars [20]. Many hospitals adopt programs using a maximum surgical blood order schedule (MSBOS). MSBOS contains a list of recommended blood orders for commonly performed surgeries regionally. A “one size fits all” solution, however, does not account for variation in patient conditions or risk beyond the surgery type. A study by Frank et al. [8] has shown that an institution-specific blood order schedule can significantly reduce unnecessary preoperative blood ordering. Our collaborator, Dr. Ryan Metcalf, has conducted statistical analysis to facilitate more accurate and precise ordering [24], but institutions have not widely adopted the novel ordering method.

3 DOMAIN GOALS

Our main collaborator, also a coauthor of this paper, is the medical director of the Transfusion Medicine Service at University of Utah. As a PBM expert, he is studying how transfusion affects patient outcomes and the overall use and effectiveness of patient blood management. Transfusion is common across specialties, including in a nonsurgical context, but we decided to focus on cardiac surgical cases because hemorrhage is more prevalent in cardiac surgeries, and cardiac surgeries are complex in nature.

To get a good understanding of the domain goals, we first interviewed Dr. Metcalf about his process for evaluating patient blood management practice. He started by comparing transfused units against preoperative hemoglobin values. A patient should first be treated for anemia if it is present, and ideally have a hemoglobin

value above 13g/dL before surgery. If a patient has a hemorrhage during the procedure, cell salvage should be considered as an alternative to transfusion. Dr. Metcalf also analyzed the patient’s post-operative hemoglobin value: a high hemoglobin value indicates excessive transfusion unless no transfusion was performed. The most important measures are patient outcomes, which can be influenced by transfusion practice. Analyzing transfusions within the context of preconditions, risk factors, details about the surgery, vital signs, and outcomes, is critical for an accurate analysis.

Dr. Metcalf also described his concerns about sharing insights effectively with surgeons, which has been only partially successful in the past. Surgeons may not always have the time to analyze data closely. When advising surgeons and anesthesiologists, he would prefer a medium that conveys the main aspects efficiently. Because a well-designed visualization can communicate the data with ease, we decided to design and implement a visual analysis tool for analyzing and reviewing patient blood management practice with outcomes of interest.

Next, to understand the needs of all stakeholders involved in PBM, we conducted a creative visualization-opportunities workshop [21]. Participants included two cardiac anesthesiologists, one cardiac surgery critical care physician, and one IT manager. The participants identified factors that potentially lead to unnecessary transfusions, and brainstormed on how to improve the current standard of care for better patient outcomes. They agreed on the need for better PBM practices, and were hopeful that a tool showing comprehensive PBM data could lead to better practices in general. The feedback from the workshop greatly helped us to establish our research plan and design for the project. After the workshop, we met with all cardiac surgeons at the University of Utah Hospital to ensure that our analysis was broadly representative.

Based on these interviews, meetings, and the creativity workshop, we identified the following goals our collaborators have regarding patient blood management:

- **Identifying problematic transfusion practices**, and thereby improving outcomes is the primary goal that all our stakeholders share. Meeting this goal requires the analysis of transfusion data in the context of patient records and evidence-based guidelines from the literature. Also, we discussed comparisons of various aspects (between clinicians, between time intervals, etc.) as a way to achieve the goal.
- **Decision support for when to transfuse** was brought up as another goal. Surgeons and anesthesiologists would like to have more data and information about when it is appropriate to transfuse while in the operating room.

When considering these goals, we quickly realized that real-time decision support would not be possible given the regulatory environment and the data we have available. Hence, we decided to focus on the retrospective analysis of transfusion practices. Based on the decided goal, we compiled a list of data items that we would need to address the analysis questions. In cooperation with the Enterprise Data Warehouse at the University of Utah, we compiled a dataset based on the electronic health records of over 4000 cardiac surgery patients, spanning the years 2014 to 2019, and covering 111 different cardiac surgery procedures, ranging from relatively benign bypass surgery to heart transplants. For each patient, we included transfusion records, medication records and lab results, for instance hemoglobin level records, thromboelastogram (ROTEM) results and units of blood products transfused, if any. We break down the high-level goal into specific design goals as follows:

- **G1: Practice-Focused Comparison.** We identified comparison as a central aspect of the analysis question. Specifically, on the highest level, our collaborators want to compare transfusion practices among (1) surgeons, (2) anesthesiologists,

and (3) years. The former two enable a comparison between individual clinicians, the latter enable temporal analysis.

- G2: Outcome and Intervention-Focused Comparison.** On top of these primary factors by which to compare cases, our collaborators want to further divide and compare cases by additional criteria, such as outcome measures (are higher numbers of transfusion associated with worse outcomes?) and time-windows (has our practice changed since the intervention implemented a year ago?)
- G3: Contextualizing.** When comparing individuals, accounting for the context of their practice is important. For example, a particular surgeon could be transfusing more than their peers because they tend to take on more complicated cases. Our interviews and workshop revealed that providing this contextual information is critical to understanding and trusting the data.
- G4: Details on Demand.** Complementary to providing contextual information in aggregate, we also need to provide details about cases on demand. To account for certain outliers, or to also understand the specific situation, clinicians need to access individual patient records.
- G5: Medical Standards and Practice.** Providing the ability to analyze practice relative to evidence-based standards is also important. For example, postoperative hemoglobin values above a threshold can indicate over-transfusion, and pre-operative hemoglobin values below a threshold can indicate a failure to manage anemia for elective (scheduled) surgeries.
- G6: Communication and Sharing.** Finally, since we have identified two types of stakeholders, our visual analysis interface needs to be tailored to both groups: PBM professionals want to interactively explore the data, create custom views, and divide based on multiple attributes, whereas clinicians want to see their performance, relative to their peers and to evidence-based standards. Hence, we envision split roles of a “power user” analyzing the data in detail, and providing annotations and context, and a “consumer” of these visualizations.

4 RELATED WORK

As digital medical records are now ubiquitous, visualizing medical records has been an important area of research for analyzing medical data. We provide an overview here, but for a more detailed survey on visualizing and exploring digital patient records, we refer to Rind et al. [28]. Caban et al. [4] describe four major types of visual analytic applications in healthcare: clinicians analyzing patients’ records, administrators making data-supported decisions, researchers working on large medical datasets, and patients understanding their own data. Our work falls into the first category. Shneiderman et al. [32] analyze the challenges of implementing interactive visualizations for healthcare professionals, which include offering busy clinicians timely information in the right format. West et al. [34] encourage creators of visualization systems to also consider the training time required. In Sanguine, we address these points by enabling a workflow where a PBM professional curates an annotated visual report for the clinicians.

In many clinical contexts, a timeline of symptoms and treatments is important; hence, multiple works have been focused on visualizing medical records in a timeline format. Lifeline [26], for example, shows details of individual patient medical histories, representing events and episodes using horizontal lines. TimeLine [3] uses a similar event-based arrangement for personal medical records, but is abundant in image and file resources for users to explore. Faiola

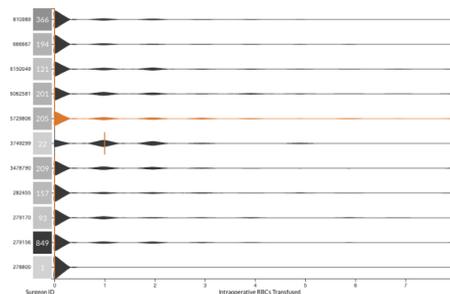


Figure 2: Our initial design used violin plots to show distributions. We abandoned this design due to the skewed and discrete distribution of the data.

et al.’s work [7] is specifically designed for ICU clinicians, juxtaposing event timelines over vital signs. DecisionFlow [16] visualizes a group of patients with similar characteristics in temporal event sequences. Outflow [35] also provides an overview of event sequences extracted from groups of patient records. All these works focus on visualizing an individual patient or a group of patients on a timeline, and emphasize specific events happening in the sequence. While event sequences can play a role in transfusions, Sanguine focuses on a provider and practice-focused approach instead.

Patient cohorts analysis is another mainstay of visualizing medical records [32]. Cohort definition and creation is a common theme in these works. COQUITO [22] uses a treemap and visual temporal queries to help users generate cohorts. Composer [29] adopts an attribute filter method for cohort creation, and compares patient cohort developments and changes over time and to other cohorts. Bernard et al. [2] also follow the attribute-oriented cohort creation, but visualize all patient records of the selected cohort, instead of comparing several cohorts through an aggregated visualization design. Phenostacks [14] uses a set visualization to compare various phenotypes over multiple cohorts simultaneously and a sunburst graph for subset and global pattern discovery. Bernard et al. [1] use static dashboards and compact data encodings to visualize groups of patients, although their approach lacks the flexibility for users to curate their own charts. Visualization studies of patient cohorts analysis typically focus on cohorts that are defined by patient characteristics, such as shared demographics. Sanguine, instead, focuses on cohorts that relate to the practice of the providers: cohorts are either created temporally, to make year-to-year comparisons, or by the provider, to compare the performance of surgeons and anesthesiologists.

Additional medical record visualizations help users navigate large datasets. Lee et al. [23] developed a web-based visualization tool to help users querying large datasets with comparative visualizations; however, the tool has limited ability for customization.

Another research topic on visualizing health record is bridging the communication gap between providers and patients [27]. As an example, PROACT [19] helps patients understand the complex risk of cancer and the treatment plans available to them through visualizations. We are not aware of existing techniques explicitly designed for provider-to-provider communication, which is a key aspect of Sanguine.

Finally, various projects have applied statistical analysis to transfusion data. A visualization tool has been developed for decision-making support on blood product ordering at a pediatric hospital [12], which gives patient-specific blood order advice. This tool allows users to filter the data based on age groups, procedures, and conditions. After filtering the data, the tool visualizes all the transfusion data and other attributes of interest in previous procedures using bar charts. This approach offers a concrete solution to a patient-specific blood order schedule, and shows how visualization can be effective for visualizing transfusion in a clinical setting. Al-

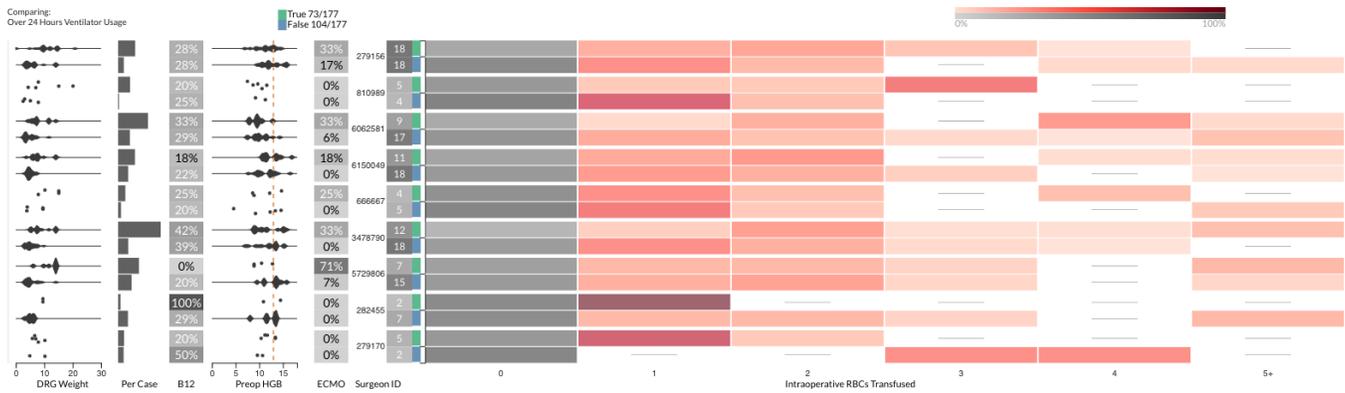


Figure 3: A heatmap showing transfusion data (red cells) for surgeons performing CABG, subdivided by the need for long-term ventilation. Cases shown in the rows with a green indicator required ventilation for more than 24 hours, whereas cases in blue rows were removed from ventilation before that, or did not require a ventilator. For context, five attributes are visualized on the left: distribution of DRG weights (risk scores), per case transfusion amount, B12 medicine usage, preoperative hemoglobin distributions, and ECMO usage rates. The color map for red blood cell transfusion is set to exclude cases without transfusions.

though the tool shows all blood usage of cases from filtered results, Sanguine offers more personalized blood ordering preparation for providers, where they can view their history of transfusions on a particular procedure with specific conditions.

5 VISUALIZATION DESIGN

The design of Sanguine reflects the goals described in Section 3, which are rooted in the interviews and creativity workshop we conducted. The subsequent development of Sanguine was done using an iterative, user-centered process based on the feedback from weekly meetings between the visualization development team and the PBM expert. Because we are designing this visualization for two user types, Sanguine needs to be powerful enough for PBM experts to discover relevant results, and be straightforward enough so that surgeons can leverage the visualization with limited time.

Sanguine consists of three main components, shown in Figure 1: a filtering and selection view, the main workspace where visualizations can be flexibly arranged, and a patient-specific detail view. A video demonstrating Sanguine is available at <https://youtu.be/DhTnyvCJgtM>. The source code is available at <https://github.com/visdesignlab/Sanguine>. Due to the sensitive nature of the data, we cannot provide a live-demo.

5.1 Data

A large amount of data is recorded in patient records, but not all of the data is typically necessary to answer a particular research question. Hence, we designed Sanguine such that a PBM expert can flexibly compose visualizations based on the data relevant to the task. The primary data that PBM experts are interested in are the number of units transfused, cell salvage (an alternative to transfusions, measured in milliliters), and laboratory measurements of hemoglobin level, which are the main indicators for anemia management and transfusion appropriateness.

Depending on the research question, other data can provide the relevant context: patient outcomes (death, need for long-term ventilation, etc.), PBM-related drug administration, and information about whether a surgery was planned or was an emergency are some example attributes. These attributes are extracted directly from electronic health records and anesthesia flowsheets.

5.2 Comparing and Contextualizing

The most important functionality of Sanguine is to visualize the transfusion data mentioned above, while addressing the comparison-related goals — comparing between practitioners and

time intervals (G1), as well as dividing and comparing using additional variables, such as outcomes (G2). Breaking these goals into abstract tasks, we need to enable the comparison of distributions of values. Following Gleicher’s guidance [13], we chose a juxtaposition strategy.

In the early development stage, we used bar charts to show the total number of transfusions, but quickly decided not to use these charts because they did not account for the number of surgeries performed, or for the number of units transfused per surgery. We then evaluated violin plots (Figure 2), since violin plots are generally well suited to visualize distributions. However, as is evident from Figure 2, the discrete nature of the data made violin plots a poor choice for visualizing these values.

Hence, we decided to use heatmaps for the main chart type for comparison visualization, as shown in Figure 1. Addressing the practice-focused comparison (G1), the heatmap shows transfusion data aggregated by surgeons, anesthesiologists, or years. To account for the difference in total case count between these facets, we use relative scales showing the percentage of cases that had 0–4 units of red blood cells, or 5 or more transfusions. The threshold varies depending on the blood component. The decision to aggregate values above a certain threshold into one bin was driven by the effect that rare outliers had on the overall visualization. For cell salvage data, which are measured in continuous values, we used bins with a dedicated bin for no usage.

As is evident from the heat map in Figure 1, a large percentage of cases received no transfusion. This aspect of the data is important to include in the visualization, but the heavily skewed distribution makes it hard to notice any difference in the nonzero transfusion cases. To address this problem, we provide a toggle to remove zero transfusion from the color-scale in the visualization — the heatmap still shows the data on zero transfusion, but shows these cases on a separate (gray) color scale, thereby making differences in the one or more transfusion cases apparent, as shown in Figure 3.

The resulting heatmap enables PBM experts and clinicians to compare their relative practice of transfusion. Although Sanguine can be used to compare transfusions across all procedures, filtering by procedures is essential, because diverse procedures such as heart transplants and bypass surgery differ in the typical need for transfusions. However, even when narrowing cases down to specific procedures, there can be systematic differences in how complicated cases are, and hence what outcomes can be expected and how much will be transfused.

To provide this context and show outcomes that are potentially related to transfusions (G3), Sanguine can visualize additional at-

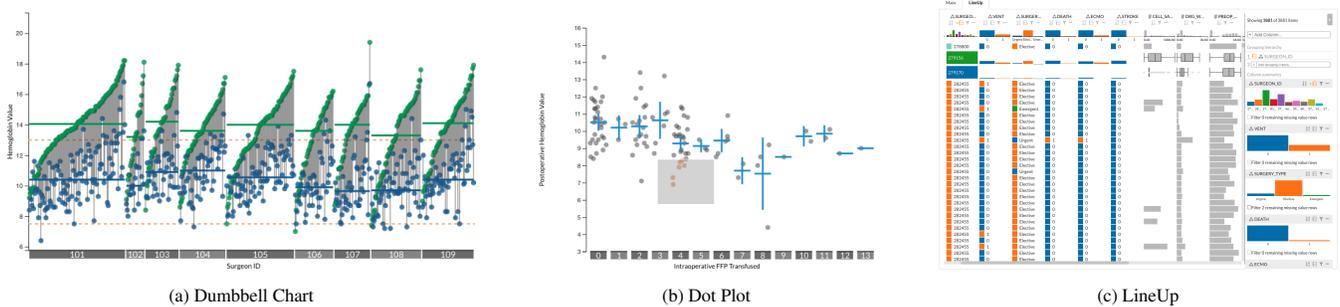


Figure 4: Visualizing bypass (CABG) procedure data: (a) A dumbbell chart that shows pre- and postoperative hemoglobin values of elective surgery cases, grouped by surgeons and sorted by preoperative hemoglobin levels. Horizontal lines in each band represent the medians, and the dotted lines are for references to preoperative hemoglobin level 13g/dL and transfusion-trigger hemoglobin level 7.5g/dL. (b) A dot plot shows the postoperative hemoglobin value to the transfused frozen plasma intraoperatively. The horizontal line in each band represents the mean, and the vertical line is the 95% confidence interval. (c) LineUp, a tabular data visualization technique, is available in Sanguine as a separate tab view. The filtering system in LineUp is connected to the main view.

tributes to display along with the heatmap, such as lab values and patient outcomes (see Figure 3). The visual encoding used depends on the attribute types. For numerical/distribution values (such as hemoglobin values), we use a violin plot that morphs into a dot plot when there are few observations; for individual numerical values, we use bar charts (e.g., average transfusions per case) or labeled heatmap cells (e.g., mortality rates).

Finally, to enable outcome and intervention-focused comparisons (G2), the heatmap in Sanguine can be divided by binary outcome variables such as mortality, or by a time-range. The heatmap shown in Figure 3 is divided by the need for long-term ventilation, as indicated by the green/blue fields. Every row is divided, showing the transfusion data and any contextual information in separate rows. In Figure 3, for example, we can see that long-term ventilation appears to correlate with higher transfusion rates and higher risk scores.

5.3 Visualization for Medical Standards and Practice

Hemoglobin laboratory values are a key indicator in patient blood management. The medical literature has established evidence-based guidelines for pre- and postoperative hemoglobin levels, and analyzing transfusion practice in this context is essential (G5). A low preoperative hemoglobin level indicates improper anemia management before surgery, whereas high postoperative hemoglobin levels imply excessive transfusion during surgery. To facilitate the analysis, Sanguine provides a dumbbell chart dedicated to hemoglobin level evaluation (see Figure 4a). We encode the recommended levels for preoperative and postoperative hemoglobin values in the chart, and they are easy to change in Sanguine through a configuration file. An individual dumbbell visualizes the preoperative (green) and the postoperative (blue) hemoglobin values of a single case as dots. A link connecting the dots shows the gap between these values. Cases and their corresponding dumbbells can be divided by all relevant attributes; Figure 4a shows dividing by surgeons. Within each division, solid horizontal lines show the medians for pre- and postoperative hemoglobin level respectively, and dotted lines provide clinically recommended values for minimal preoperative hemoglobin level and transfusion-trigger hemoglobin level. The dumbbells within each division can be sorted based on preoperative value, postoperative value, or the gap between the two.

5.4 Views for Visualization, Filtering, and Brushing

In addition to dynamic comparisons, filters and brushes are essential tools to create the visualization relevant to answer specific analysis questions. In addition to the procedure filter view (Figure 1), Sanguine provides views specifically designed for filtering based on data values. The dot plot (Figure 4b) can visualize correlations

and support rectangle brushes, which can be converted into filters in the filter manager view (Figure 1). The dot plot also visualizes mean values and confidence intervals.

We also integrate a LineUp [11, 17] view (Figure 4c) to visualize all attributes of individual cases (G4) and complement Sanguine’s filter system with the one that comes with LineUp. LineUp visualizes data in a tabular layout, and applies different techniques to columns based on the attribute types.

5.5 Case View

An alternative way to view information about individual cases (G4) is the detail view, shown in Figure 1 on the right. All selected cases are shown in a list; the attributes for one case are shown, including transfusion record, surgery description, relevant medicines administered, and outcomes. Using the detail view, analysts can study these cases closely, and if they see something interesting, they can add that attribute to another visualization in Sanguine. Ultimately, practitioners can also cross-reference cases with the medical record system, which contains even more data on the patients.

5.6 Communication and Sharing

Sanguine enables communication and sharing (G6) in two complementary ways: first, communication is facilitated through annotations: each visualization is accompanied with a text field, which can be used for notes or to record insights and conclusions.

Second, sharing of findings is enabled through provenance tracking. Each change to the visualization, including visualization configurations, filter settings, and annotations, is saved as a state, and recorded using the Ttrack provenance tracking library [5]. Based on Ttrack’s functionality, Sanguine provides undo/redo, saving and loading the state of a workspace to/from a server, and sharing the state via a URL. URL based sharing is convenient for distributing findings made by the PBM expert to clinicians, who can then review the visualizations and adjust their practice, if appropriate. As these “visualization consumers” typically do not want to leverage the full complexity of Sanguine, we introduced a “View Mode”, which removes editing functionality and simplifies the interface.

5.7 Implementation and Deployment

Sanguine is developed in Typescript using the React framework and D3, supported by a Django server, and interfaces with an SQL database housed in the Enterprise Data Warehouse. It is deployed in a protected environment suitable for sensitive medical data, and uses encryption and two-factor authentication to ensure security.

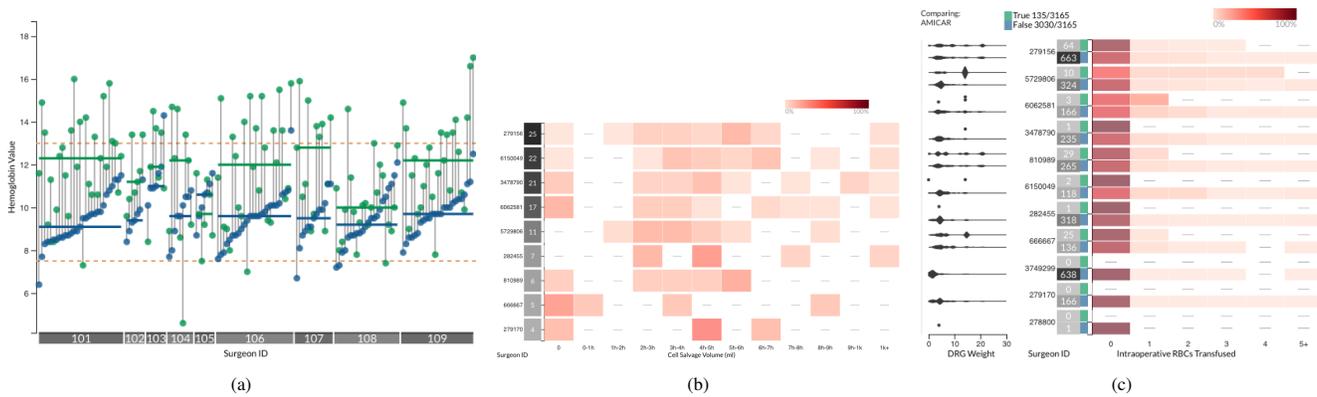


Figure 5: Case studies. (a) A dumbbell chart comparing surgeons and sorted by postoperative hemoglobin levels. Cases shown are with one or more units of red blood cell transfusions. The high postoperative hemoglobin values (above the lower dashed line) indicate that over-transfusion is widespread. (b) A heatmap of cell salvage usage, for cases with one or more units of red blood cells transfused. Several surgeons do not or rarely use cell salvage when transfusing, which violates best practices. (c) A heatmap of red blood cell transfusion, divided by the use of aminocaproic acid, a drug that can reduce bleeding. Many surgeons do not use the drug. The violin plot shows DRG weights (risk scores). Higher risk patients tend to receive aminocaproic acid more often.

6 CASE STUDY

To demonstrate the utility of Sanguine, we present case studies of two scenarios where a PBM expert, Dr. Metcalf, uses Sanguine (1) for analysis of transfusion practices, and (2) as a “patients like mine” decision support tool when ordering blood components during preparation for surgery.

6.1 PBM Review and Analysis

Anemia Management. Preoperative anemia management is an essential part of patient blood management. Dr. Metcalf started his analysis by studying whether patients who had anemia were treated for it before surgery, as anemia can require transfusions that would otherwise not be necessary. A patient with a hemoglobin level at 10g/dL or below is considered to be anemic. To view the overall hemoglobin trend, Dr. Metcalf first created a dumbbell chart comparing surgeons, shown in Figure 4a. Since preoperative anemia management can be done only for elective surgeries, he then used the surgery urgency filter in the filter manager to remove emergent and urgent cases. He also selected a common procedure CABG (coronary artery bypass grafting) in the filter selection, and analyzed a subset of surgeons who have performed CABG. By comparing preoperative hemoglobin values to the dotted line for 13g/dL as the clinically recommended preoperative hemoglobin values, he noticed several cases should have been treated for anemia before surgery, but were not. He also noticed that some surgeons seemed to be more careful about anemia management than others.

Transfusion Appropriateness. Using the same dumbbell chart, Dr. Metcalf also reviewed transfusion appropriateness. When a patient receives a red blood cell transfusion, the provider’s target postoperative hemoglobin value should be between 7g/dL and 9g/dL. To view only patients who received red blood cell transfusions, Dr. Metcalf added a dot plot of postoperative hemoglobin and units of red blood cells transfused intraoperatively. Appropriate transfusion is a practice applicable to all cases, elective or urgent; hence, he removed the elective case filter, and changed the sorting of the dumbbell chart to postoperative hemoglobin levels for easier comparison of the postoperative values. Using the brush feature of the dot plot, Dr. Metcalf could filter out patients who did not receive red blood cell transfusions during the surgery, resulting in the dumbbell chart shown in Figure 5a. After applying the filter, Dr. Metcalf observed a significant amount of cases for which postoperative hemoglobin values were much higher than the target value,

indicating over-transfusion in these cases. Dr. Metcalf then used this information to advise surgeons and anesthesiologists for better hemoglobin level targeting when performing transfusions.

Cell Salvage. Another key point of patient blood management is the use of cell salvage. Cell salvage is recycling a patient’s blood when they bleed during the surgery, an alternative to transfusions. Ideally, all surgeons should use cell salvage when performing transfusions. Dr. Metcalf started with a heatmap of cell salvage, which he aggregated by surgeons. The heatmap shows that most surgeons used cell salvage, but the chart does not indicate how surgeons were using cell salvage when they were also transfusing.

After removing all cases without transfusions, Dr. Metcalf observed that even though cell salvage was used most cases, there was still room for improvement, as shown in Figure 5b. Most providers were using cell salvage properly; however, a few providers were performing transfusions without using cell salvage for over 20% of their cases. This information is valuable for PBM experts, who can now identify the providers and inform them about properly using cell salvage.

Drug Treatment. Finally, Dr. Metcalf analyzed usage of antifibrinolytic agents, in this case, aminocaproic acid. Antifibrinolytic agents help to reduce bleeding in surgeries, and overall reduce the use of transfusion. To visualize the use of aminocaproic acid, Dr. Metcalf added a comparison heatmap of red blood cells transfused, aggregated by surgeon, divided by aminocaproic acid usage, shown in Figure 5c. The chart gave him an idea of how frequently providers were administering aminocaproic acid, and how effective it was at reducing transfusions. By analyzing the heatmap and comparing the case counts, he concluded that the majority of cases did not receive aminocaproic acid, which is not optimal for patient blood management in cardiac surgeries.

Dr. Metcalf could not find a consistent trend of fewer transfusions associated with the use of aminocaproic acid, possibly due to the low absolute number of cases that utilized it. To investigate the use of aminocaproic acid in more context, he added plots of the distribution of the DRG weights (risk scores). He noticed that cases that used aminocaproic acid but still received several transfusions tended to have higher DRG weights, which indicates that these patients were at higher risk for surgical complications.

Dr. Metcalf remarked that none of these analyses were possible before, at least not with the flexibility that Sanguine enables.

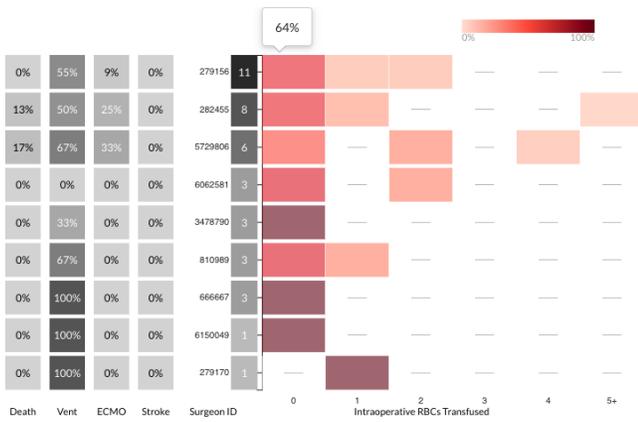


Figure 6: Heatmap of units of red blood cells transfused for the open MVR procedure and cases with preoperative hemoglobin levels of 10-12g/dL. By narrowing the cases down to procedures and parameters specific to the case, surgeons can plan how many units of blood to order.

6.2 “Patients Like Mine” Decision Support

Although Sanguine is designed primarily for retrospective PBM analysis, Dr. Metcalf also explored using Sanguine as a decision support tool for blood component ordering when preparing for surgeries. For example, a surgeon preparing for an open mitral valve replacement (MVR) can view the transfusion data for prior open MVR cases, using Sanguine’s heatmap shown in Figure 6. To get a more specific picture, they can apply a filter using the dot plot based on their patient’s preoperative hemoglobin value, and even filter based on the anesthesiologist with whom they will work in the surgery. The surgeon can view their own historic transfusion for open MVR, and order blood components based on their personal records. For example, a surgeon who performed 11 open MVR cases (first row in Figure 6) did not transfuse in 64% of cases, and used only 1 unit in 15% of all their cases. Hence, they can conclude that they should order one unit of red blood cells before surgery.

Surgeons can also view the outcomes of cases from historic data, and know what they can expect for their upcoming case. For example, for the same surgeon in the first row, they can see that 55% of the cases required long-term ventilation after the surgery. Overall, Dr. Metcalf remarked that using Sanguine is a much more personalized and precise way for surgeons to prepare, compared to the guidelines currently in place. He commented that in this way, Sanguine can help reduce waste, while at the same time also reducing the need for emergency blood releases during the surgery.

7 FEEDBACK FROM CLINICIANS

Dr. Metcalf, the PBM expert, presented his findings, described in the previous section, to a cardiac surgeon in a virtual meeting, using a screen-share of Sanguine in “View Mode”. The interface and visualization techniques were well received, and the surgeon was positive that Sanguine can be an effective tool to study transfusions in their clinical context. They said, *this is what I have been looking into and could not find an answer anywhere. Besides just blood transfusion, there is a lot of other data in there too...*

We asked them to compare Sanguine with the current PBM analysis tool available for surgeons. The current practice for cardiac surgeons requires them reporting to a national database of The Society of Thoracic Surgeons (STS), and each hospital cardiac department receives a semi-annual report. As a result, the report is delayed, and not effective for a timely analysis of PBM practice. In contrast, Sanguine is directly connected to the hospital database, which can be updated frequently. The surgeon commented on the difference:

the STS reports we currently receive are about six months delayed, and you can update this[Sanguine] much faster...This can be used for meetings to see where we can improve.

Speaking to a weakness, the surgeon raised a concern regarding the integrity of some of the data. First, they questioned the validity of the cell salvage values, because sometimes the usage is not properly captured in the anesthesia flowsheets, from which Sanguine derives the data. They also were surprised by and doubted the accuracy of the minimal usage of aminocaproic acid, which they expected everyone to use on every case for cardiac surgery. The sources for these errors are unknown, but they are present in the data that Sanguine is built on, and this is not a limitation of Sanguine itself.

8 DISCUSSION AND LIMITATION

Evaluation. Even though Sanguine was well received by clinicians and PBM experts during our development process, in this paper, we validate Sanguine through limited case studies with our collaborator and one surgeon. As a next step, we want to see if Sanguine can be adopted in the regular routine of the hospital PBM review process without a long learning curve. We plan on collecting longitudinal data on how Sanguine is adopted as a regular analytic tool and adapt our design based on the feedback, yet this long-term evaluation is beyond the scope of this paper.

Data Integrity. As described in Section 7, the integrity of some of the data was questioned when we presented Sanguine to clinicians. We have yet to determine the source of the problem, but a tool like Sanguine can make the clinical team aware of deficiencies in the data collection process. However, this problem indicated that it will be important for Sanguine to provide a better annotation system, enabling PBM experts and clinicians to pinpoint problems.

Data Integration. The dataset we used to develop Sanguine is static and deidentified; however, the deidentification procedures are not certified and, hence, we cannot show certain data fields in this paper. Since we query the data directly from the Enterprise Data Warehouse through SQL queries, and the data types were not altered during deidentification, we expect a smooth transition from deidentified dataset to the original dataset for clinical use.

9 CONCLUSION AND FUTURE WORK

In this paper, we introduced Sanguine, a tool for analyzing patient blood management data. Sanguine is currently deployed and our collaborator is in the process of adopting Sanguine in his regular workflow. Going forward, we plan to evaluate the long-term use of Sanguine and also study the influence it has on the decision-making process. We also plan to integrate several additional datasets, such as ROTEM results and platelet counts. We aim to make Sanguine reports available to surgeons in the near future.

We would also like to explore expanding the Sanguine annotation system, as described in Section 8. A system that allows analysts to annotate on any components would enhance our current text-based chart annotation structure, and further facilitate communication between PBM experts and clinicians.

Even though we focus on PBM data in this design study, the techniques we used can likely be applied to other areas of medicine, where analysts want to quickly view outcomes of interest and their relationships with variables of interest.

ACKNOWLEDGMENTS

We thank Dr. Vikas Sharma, the Enterprise Data Warehouse, and the Center for High Performance Computing at the University of Utah. We gratefully acknowledge funding for this project by ARUP Laboratories. Computational resources used were partially funded by the NIH Shared Instrumentation Grant 1S10OD021644-01A1.

REFERENCES

- [1] J. Bernard, D. Sessler, J. Kohlhammer, and R. A. Ruddle. Using dashboard networks to visualize multiple patient histories: A design study on post-operative prostate cancer. *IEEE transactions on visualization and computer graphics*, 25(3):1615–1628, 2018.
- [2] J. Bernard, D. Sessler, T. May, T. Schlomm, D. Pehrke, and J. Kohlhammer. A Visual-Interactive System for Prostate Cancer Cohort Analysis. *IEEE Computer Graphics and Applications*, 35(3):44–55, May 2015. doi: 10.1109/MCG.2015.49
- [3] A. A. T. Bui, D. R. Aberle, and H. Kangaroo. TimeLine: Visualizing Integrated Patient Records. *IEEE Transactions on Information Technology in Biomedicine*, 11(4):462–473, 2007. doi: 10.1109/TITB.2006.884365
- [4] J. J. Caban and D. Gotz. Visual analytics in healthcare - opportunities and research challenges. *Journal of the American Medical Informatics Association*, 22(2):260–262, Mar. 2015. doi: 10.1093/jamia/ocv006
- [5] Z. T. Cutler, K. Gadhave, and A. Lex. Trrack: A Library for Provenance Tracking in Web-Based Visualizations. Preprint, Open Science Framework, July 2020. doi: 10.31219/osf.io/wncbt
- [6] K. D. Ellingson, M. R. P. Sapiano, K. A. Haass, A. A. Savinkina, M. L. Baker, K.-W. Chung, R. A. Henry, J. J. Berger, M. J. Kuehnert, and S. V. Basavaraju. Continued decline in blood collection and transfusion in the United States—2015. *Transfusion*, 57(Suppl 2):1588–1598, June 2017. doi: 10.1111/trf.14165
- [7] A. Faiola. Medical information visualization assistant system and method, Feb. 2014.
- [8] S. M. Frank, M. J. Oleyar, P. M. Ness, and A. A. R. Tobian. Reducing Unnecessary Preoperative Blood Orders and Costs by Implementing an Updated Institution-specific Maximum Surgical Blood Order Schedule and a Remote Electronic Blood Release System. *Anesthesiology*, 121(3):501–509, 2014. doi: 10.1097/ALN.0000000000000338
- [9] S. M. Frank, L. M. Resar, J. A. Rothschild, E. A. Dackiw, W. J. Savage, and P. M. Ness. A novel method of data analysis for utilization of red blood cell transfusion: Analysis of RBC Utilization. *Transfusion*, 53(12):3052–3059, Dec. 2013. doi: 10.1111/trf.12227
- [10] S. M. Frank, W. J. Savage, J. A. Rothschild, R. J. Rivers, P. M. Ness, S. L. Paul, and J. A. Ulatowski. Variability in Blood and Blood Component Utilization as Assessed by an Anesthesia Information Management System. *Anesthesiology*, 117(1):99–106, July 2012. doi: 10.1097/ALN.0b013e318255e550
- [11] K. Furmanova, S. Gratzl, H. Stitz, T. Zichner, M. Jaresova, A. Lex, and M. Streit. Taggle: Combining overview and details in tabular data visualizations. *Information Visualization*, 19(2):114–136, 2020. doi: 10.1177/1473871619878085
- [12] J. A. Gálvez, L. Ahumada, A. F. Simpaio, E. E. Lin, C. P. Bonafide, D. Choudhry, W. R. England, A. F. Jawad, D. Friedman, D. A. Sesok-Pizzini, and M. A. Rehman. Visual analytical tool for evaluation of 10-year perioperative transfusion practice at a children’s hospital. *Journal of the American Medical Informatics Association*, 21(3):529–534, May 2014. doi: 10.1136/amiajnl-2013-002241
- [13] M. Gleicher. Considerations for Visualizing Comparison. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):413–423, Jan. 2018. doi: 10.1109/TVCG.2017.2744199
- [14] M. Glueck, A. Gvozdiak, F. Chevalier, A. Khan, M. Brudno, and D. Wigdor. PhenoStacks: Cross-Sectional Cohort Phenotype Comparison Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):191–200, 2017. doi: 10.1109/TVCG.2016.2598469
- [15] L. T. Goodnough and A. Shander. Patient blood management. *Anesthesiology*, 116(6):1367–1376, June 2012. doi: 10.1097/ALN.0b013e318254d1a3
- [16] D. Gotz and H. Stavropoulos. DecisionFlow: Visual Analytics for High-Dimensional Temporal Event Sequence Data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1783–1792, Dec. 2014. doi: 10.1109/TVCG.2014.2346682
- [17] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. LineUp: Visual Analysis of Multi-Attribute Rankings. *IEEE Transactions on Visualization and Computer Graphics (InfoVis '13)*, 19(12):2277–2286, 2013. doi: 10.1109/TVCG.2013.173
- [18] I. Gross, B. Seifert, A. Hofmann, and D. R. Spahn. Patient blood management in cardiac surgery results in fewer transfusions and better outcome. *Transfusion*, 55(5):1075–1081, 2015. doi: 10.1111/trf.12946
- [19] A. Hakone, L. Harrison, A. Ottley, N. Winters, C. Gutheil, P. K. J. Han, and R. Chang. PROACT: Iterative Design of a Patient-Centered Visualization for Effective Prostate Cancer Health Risk Communication. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):601–610, Jan. 2017. doi: 10.1109/TVCG.2016.2598588
- [20] T. Hannon. Waste Not, Want Not. *American Journal of Clinical Pathology*, 143(3):318–319, 2015. doi: 10.1309/AJCPM8FACVC0RPRG
- [21] E. Kerzner, S. Goodwin, J. Dykes, S. Jones, and M. Meyer. A Framework for Creative Visualization-Opportunities Workshops. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):748–758, Jan. 2019. doi: 10.1109/TVCG.2018.2865241
- [22] J. Krause, A. Perer, and H. Stavropoulos. Supporting Iterative Cohort Construction with Visual Temporal Queries. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):91–100, Jan. 2016. doi: 10.1109/TVCG.2015.2467622
- [23] J. Lee, E. Ribey, and J. R. Wallace. A web-based data visualization tool for the MIMIC-II database. *BMC Medical Informatics and Decision Making*, 16(1):15, Feb. 2016. doi: 10.1186/s12911-016-0256-9
- [24] R. A. Metcalf, M. B. Pagano, J. R. Hess, J. Reyes, J. D. Perkins, and M. I. Montenovio. A data-driven patient blood management strategy in liver transplantation. *Vox Sanguinis*, 2018. doi: 10.1111/vox.12650
- [25] A. Pfuntner, L. M. Wier, and C. Stocks. Most Frequent Procedures Performed in U.S. Hospitals, 2010: Statistical Brief #149. In *Healthcare Cost and Utilization Project (HCUP) Statistical Briefs*. Agency for Healthcare Research and Quality (US), Rockville (MD), 2013.
- [26] 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 B. B. Bederson and B. Shneiderman, eds., *The Craft of Information Visualization*, Interactive Technologies, pp. 308–312. Morgan Kaufmann, San Francisco, Jan. 2003. doi: 10.1016/B978-155860915-0/50038-X
- [27] Y. G. Rajwan and G. R. Kim. Medical information visualization conceptual model for patient-physician health communication. In *Proceedings of the 1st ACM International Health Informatics Symposium, IHI '10*, pp. 512–516. Association for Computing Machinery, New York, NY, USA, Nov. 2010. doi: 10.1145/1882992.1883074
- [28] A. Rind. Interactive Information Visualization to Explore and Query Electronic Health Records. *Foundations and Trends® in Human-Computer Interaction*, 5(3):207–298, 2013. doi: 10.1561/1100000039
- [29] J. Rogers, N. Spina, A. Neese, R. Hess, D. Brodke, and A. Lex. Composer—Visual Cohort Analysis of Patient Outcomes. *Applied Clinical Informatics*, 10(02):278–285, 2019. doi: 10.1055/s-0039-1687862
- [30] S. Sahu, Hemlata, and A. Verma. Adverse events related to blood transfusion. *Indian Journal of Anaesthesia*, 58(5):543–551, 2014. doi: 10.4103/0019-5049.144650
- [31] A. Shander, M. Javidroozi, S. Perelman, T. Puzio, and G. Lobel. From Bloodless Surgery to Patient Blood Management. *Mount Sinai Journal of Medicine: A Journal of Translational and Personalized Medicine*, 79(1):56–65, Jan. 2012. doi: 10.1002/msj.21290
- [32] B. Shneiderman, C. Plaisant, and B. W. Hesse. Improving Healthcare with Interactive Visualization. *Computer*, 46(5):58–66, May 2013. doi: 10.1109/MC.2013.38
- [33] The Joint Commission and the American Medical Association- Convened Physician Consortium for Performance Improvement. Proceedings from the National Summit on Overuse, 2012.
- [34] 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, Mar. 2015. doi: 10.1136/amiajnl-2014-002955
- [35] K. Wongsuphasawat and D. Gotz. Exploring Flow, Factors, and Outcomes of Temporal Event Sequences with the Outflow Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2659–2668, Dec. 2012. doi: 10.1109/TVCG.2012.225