Proceedings of the 2016

Workshop on Visual Analytics in Healthcare:

November 12th, 2016
Chicago, IL
www.visualanalyticshealthcare.org

Agenda:

Session I: Introduction to Visual Analytics in Healthcare Chair: Bhanu Bahl, PhD Harvard Medical School	
8:30 - 8:45	Welcome – VAHC 2016
8:45 - 9:30	*Introduction to Visual Analytics in Healthcare* Jeremy Warner, MD Vanderbilt University
9:30 - 10:00	**Paper presentations**
	"My List vs. the Other List: The Design of an Interactive Visual Interface for Medication Reconciliation" Yarden Livnat, Bryan Gibson, Heidi Kramer, Ab Brody, Iona Thraen and Randall Rupper
	"Network analysis of treatment patterns in breast cancer care" Shane Weisberg, Douglas Hill, Rebecca Faill and Amar Das
10:00 - 10:30	Coffee Break

Session II: Design Principles Chair: Danny Wu, PhD MSI University of Cincinnati	
10:30 - 11:15	*Data Visualization - Design Principles* Shira Fischer, MD, PhD RAND Corporation
11:15 - 12:00	**Paper presentations**
	"m-TSNE: A Framework for Visualizing High-Dimensional Multivariate Time Series" Minh Nguyen, Sanjay Purushotham, Hien To and Cyrus Shahabi
	"*Interactive Visualization and Exploration of Patient Progression in a Hospital Setting"* Wathsala Widanagamaachchi, Yarden Livnat, Peer-Timo Bremer, Scott Duvall and Valerio Pascucci
	"Understanding Care Plans of Community Acquired Pneumonia Based on Sankey Diagram" Shunan Guo, Chaoguang Lin, David Gotz, Bo Jin, Hongyuan Zha, Linhua Shu and Nan Cao

Session III: Poster and Demonstrations Chair: Jesus J Caban, PhD Walter Reed National Military Medical Center	
1:00 – 1:30	Fast forward Presentations – Posters and Live Demonstrations
1:30 – 3:00	Poster and Demo Session *(see next few pages for list of posters & live demonstrations)*

Session IV: Validation and Design Chair: Uba Backonja, PhD RN University of Washington School of Medicine	
3:00-3:45	*Results - VAHC Mini Design Challenge - How? Why?*
3:00 – 4:30	**Paper Presentations**
	"Visualization for All: The Importance of Creating Data Representations Patients Can Use" Carolyn Petersen
	"Clinically Relevant Filters to Consider When Designing A Visualization for Longitudinal Electronic Health Records" Filip Dabek, Jefferson McMillan and Jesus Caban
	"Evaluating Visual Analytics for Health Informatics Applications: A Progress Report from the AMIA" David Gotz, David Borland, Jesus Caban, Dawn Dowding, Brian Fisher, Vadim Kagan and Danny T.Y. Wu
4:30 - 4:35	Closing Remarks

Poster Presentations (1:00 – 3:00pm)	
Poster #1	*"Visualizing Edit Distances Between Kinase Inhibitor Names and English Words"* Sandeep Jain and Jeremy Warner
Poster #2	*"CHAI: A Visual Interface for Examining Subject Matter Similarities across Intervention Chat Message Histories"* Annie T. Chen, William R. Kearns, Emily F. Law, Nicole M. Alberts and Tonya M. Palermo
Poster #3	*"An Interactive Tool for Clinical Data Mining and Visualization"* Qianxi Li, Yugang Jia, Merlijin Sevenster, Laura Olivieri, Craig Sable, Yue-Hin Loke and Justin Wiggs
Poster #4	*"Genetics-based Motivation Viz ("GMV"): Visualizing Direct-to-Consumer Genetic Test Results To Empower Health Behavior Change"* Kenyon Crowley, Joohee Choi and Rohan Bondili
Poster #5	*"Rapid, Iterative Design of reCAP, the REal-Time Care Analysis Platform, to Support Care Redesign"* Michael Ripperger and Colin Walsh
Poster #6	*"Visualizing Patient Genomic Data and Predicted Drug Sensitivities in the SMART Precision Cancer Medicine Application"* Krysten Harvey and Jeremy Warner
Poster #7	*"Visualization of EHR Data for Decision-Making in Diabetes and Congestive Heart Failure"* Shira Fischer, Charles Safran, Krzysztof Gajos and Adam Wright
Poster #8	*"Visualizing State-Level Chronic Disease Indicators as a Prelude to Insight on Meeting Healthy People 2020 Objectives"* Umesh Singh and Victoria Wangia-Anderson

Live Demonstrations (1:00 – 3:00pm)	
Demo #1	*"Evolving Visual Analytics for Better Clinical Decisions"* Dave Anderson
Demo #2	*"A semantic search engine for integration and visualization of Electronic Health Record data and enrichment with open access and public life sciences data sources"* Filip Pattyn, Bhanu Bahl and Hans Constandt
Demo #3	*"A Clinical Data Analytics Workbench to Streamline Analytics Tasks and Visualize Key Results"* Yiqin Yu, Xiang Li, Haifeng Liu, Bibo Hao, Wen Sun and Guotong Xie
Demo #4	*"Peeking into Patterns of Clinical Event Sequences with Peekquence"* Bum Chul Kwon, Janu Verma and Adam Perer
Demo #5	*"Using Advanced Analytics and Visualization to Detect HIPAA Violations"* Nick Culbertson and Robert Lord
Demo #6	*"Visualizing State-Level Chronic Disease Indicators as a Prelude to Insight on Meeting Healthy People 2020 Objectives"* Umesh Singh and Victoria Wangia-Anderson
Demo #7	*"Visualizing Patient Genomic Data and Predicted Drug Sensitivities in the SMART Precision Cancer Medicine Application"* Krysten Harvey and Jeremy Warner
Demo #8	*"Genetics-based Motivation Viz ("GMV"): Visualizing Direct-to-Consumer Genetic Test Results To Empower Health Behavior Chang"* Kenyon Crowley, Joohee Choi and Rohan Bondili
Demo #9	*"Rapid, Iterative Design of reCAP, the REal-Time Care Analysis Platform, to Support Care Redesign"* Michael Ripperger and Colin Walsh
Demo #10	*"Visual Summarization of Longitudinal Clinical Trajectories to Improve Population Health Analysis"* Filip Dabek and Jesus J Caban

Proceedings of the 2016
Workshop on Visual Analytics in Healthcare

PAPER PRESENTATIONS

My List vs. the Other List: The Design of an Interactive Visual Interface for Medication Reconciliation

Y. Livnat, B. S. Gibson, H. S. Kramer, A. A. Brody, I. Thraen, R. Rupper

***Abstract*—We present an interactive visual interface for medication reconciliation intended to provide cognitive support for VA physicians in the context of home health. We describe the iterative design process, the software prototype implementation and present insights from our user study. In particular, we provide a detailed discussions of key design decisions, the principles behind them and their implications. Though home health setting introduces some unique challenges, we believe that the reasoning behind our design decisions and the insights gained through this work can benefit the design of future medication reconciliation tools for use in other settings.**

***Index Terms*—Medication Reconciliation, Visualization, Design, Home Health, Transitional Care.**

I. INTRODUCTION

Patients transitioning from one healthcare setting to another are at high risk for adverse events and readmission [1][2][3]. One cause of these adverse events is a failure to correctly reconcile medications across multiple plans of care [4][5][6]. Medication reconciliation is a complex task that is time consuming, laborious and prone to human errors. A provider engaged in the process of medication reconciliation must detect differences between medication lists, determine the reasons for the discrepancies, decide which prescription is most appropriate for the patient, and finally change medication orders and update the patient records [7]. Additional complications arise because the medication lists are collected from disparate sources, are often long and complex, and the differences between them are often subtle. An efficient and accurate medication reconciliation process has become an important patient safety concern [8].

In this project we developed a system for medication reconciliation performed by a provider whose patient was refereed for home healthcare. The process of medication reconciliation in home healthcare is different from other settings in which the provider works directly with the patient. In a home health (HH) setting, a home health nurse visits the patient's home and reports about medications found there and patient adherence to the prescribed regimen to the certifying physician. The communication, mediated by the nurse, between the referring provider and the patient is therefore indirect and asynchronous and in some cases requires several iterations in order to verify information and correct the patient medications [8][9].

To date, technology based solutions to medication reconciliation have been mostly directed towards hospital settings and automated systems that streamline the process by gathering medication lists from disparate sources and identifying matches and discrepancies between them [2][5][6]. In contrast, very little has been published on interactive visualization systems that focus on the clinical decision-making during the discrepancy resolution phase.

In this paper, we present an interactive visual interface for medication reconciliation for physicians treating patients receiving home health services that aims to provide cognitive support for both discrepancy detection and resolution. We describe the iterative design process, the software prototype implementation and present insights from our user study [11].

The main contributions of this work are in the approach that we codify as "my list vs. the other list", the detailed discussions of key design decisions, the principles behind them and their implications. We note that while the home health setting introduces unique challenges, the overall design and the reasoning behind our design decisions are applicable in other settings.

II. HOME HEALTH USE CASE

In this work we focus on the use case of medication reconciliation specific to HH care. After a patient is discharged from a hospital, a referral for HH is generated by a physician or their designee, signed electronically by the physician and then sent to a HH agency. If the referral includes skilled nursing care, a registered nurse is sent to the

This work was supported in part by the U.S. Department of Veterans Affairs Office of Rural Health under Grant N19-FY14Q1-S3-P01023.

Y Livnat is with the Scientific Computing and Imaging Institute at the University of Utah, UT (email: yarden@sci.utah.edu)

B. S. Gibson is with the Department of Biomedical Informatics at the University of Utah School of Medicine, UT (email: bryan.gibson@utah.edu)

H. S. Kramer is with the Salt Lake VA, UT and the Department of Biomedical Informatics at the University of Utah School of Medicine, UT (email: heidi.kramer@va.gov)

A. A. Brody is with the Bronx VA Geriatric Research and Education Center and New York University Rory Meyers College of Nursing, NY (email; ab.brody@nyu.edu)

I. Thraen is with the Utah Department of Health and the University of Utah College of Social Work, UT (email: ithraen@utah.gov)

R. Rupper is with the Salt Lake VA Geriatric Research Education and Clinical Center, UT (email: Randall.Rupper@va.gov)

patient's home to conduct an initial assessment, develop a plan of care and initiate the plan of care. During the initial assessment the nurse reviews the medications found at the home and notes any differences between the patient reported medication regimen and the one listed on the referral. The home health plan of care is then reported in paper form (aka CMS-485) back to the referring providers for reconciliation and approval.

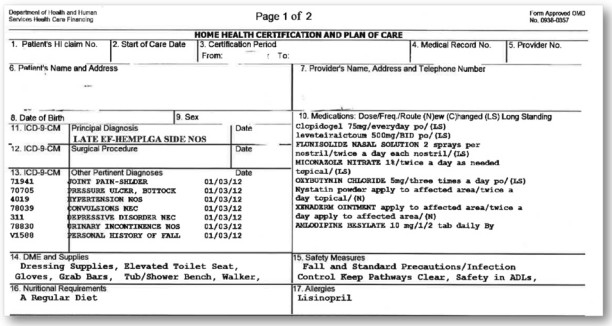

Department of Health and Human Services Health Care Financing — Page 1 of 2 — Form Approved OMB No. 0938-0357

HOME HEALTH CERTIFICATION AND PLAN OF CARE

1. Patient's HI claim No.	2. Start of Care Date	3. Certification Period From: To:	4. Medical Record No.	5. Provider No.

6. Patient's Name and Address

7. Provider's Name, Address and Telephone Number

8. Date of Birth | 9. Sex

11. ICD-9-CM Principal Diagnosis: LATE EF-HEMPLGA SIDE NOS — Date

12. ICD-9-CM Surgical Procedure — Date

13. ICD-9-CM	Other Pertinent Diagnoses	Date
71941	JOINT PAIN-SHLDER	01/03/12
70705	PRESSURE ULCER, BUTTOCK	01/03/12
4019	HYPERTENSION NOS	01/03/12
78039	CONVULSIONS NEC	01/03/12
311	DEPRESSIVE DISORDER NEC	01/03/12
78830	URINARY INCONTINENCE NOS	01/03/12
V1588	PERSONAL HISTORY OF FALL	01/03/12

10. Medications: Dose/Freq./Route (N)ew (C)hanged (LS) Long Standing

Clopidogel 75mg/everyday po/(LS)
levetairetzum 500mg/BID po/(LS)
FLUNISOLIDE NASAL SOLUTION 2 sprays per nostril/twice a day each nostril/(LS)
MICONAZOLE NITRATE 1%/twice a day as needed topical/(LS)
OXYBUTYNIN CHLORIDE 5mg/three times a day po/(LS)
Nystatin powder apply to affected area/twice a day topical/(N)
XRADEXM OINTMENT apply to affected area/twice a day apply to affected area/(N)
AMLODIPINE BESYLATE 10 mg/1/2 tab daily By

14. DME and Supplies
Dressing Supplies, Elevated Toilet Seat, Gloves, Grab Bars, Tub/Shower Bench, Walker,

15. Safety Measures
Fall and Standard Precautions/Infection Control Keep Pathways Clear, Safety in ADLs,

16. Nutritional Requirements
A Regular Diet

17. Allergies
Lisinopril

Figure 1 A portion of a CMS-485 form. The home health medications are listed in box 10 on the right.

III. Related Work

Medication errors and adverse events related to medications are increasingly recognized as significant causes of harm to patients [2][5][6] and that an accurate medication reconciliation is critical for patient safety . The Joint Commission on Accreditation of Healthcare Organizations (JCAHO) has identified [12] five core steps in medication reconciliation: *gather* information about current medications, *develop* a list of medications to be prescribed, *identify* discrepancies, make clinical decisions to *resolve* discrepancies, and *communicate* the new list. Efforts have been mostly focused on automating the *gathering*, *identification* and *communication* steps but there is very little published work on visual interfaces to support the clinical decision making process[13].

One example of a visual interface to support the medication review is the Pre-Admission Medication List (PAML)[14]. The interface compiles all medications into a single list and provides a centralized view of all currently available information about the medications. The system focuses on providing comprehensive information but it does not directly provide cognitive support in the sense of contrasting and highlighting discrepancies or facilitating appropriate actions based on the decisions the user makes.

Recently, Plaisant et al. [15], [16] presented TwinList, a user interface for reconciling two medication lists in a hospital setting that aims directly at facilitating the decision making process. The main ideas behind TwinList are the separation of the two medication lists (A and B) into five lists representing different types of discrepancies, the spatial layout of these lists in the display, and the use of a multi-step animation to illustrate the creation of the discrepancy lists. The five discrepancy lists in TwinList represent: 1) exact match, 2) medications in list A that are similar to but do not have an exact match in B, 3) medications in list B similar to A but not exact, 4) medications unique to A, and 5) medication unique to B. The five lists characterize the possible relationships between the two medication lists. In addition, animation plays a key role in TwinList and the authors suggest it may help users learn and understand the layout.

The TwinList visualization of the five lists pose a few challenges. First, the user must contend with comprehending and contrasting five lists rather than only two. Second, the exact-match list (medications with no discrepancies) is placed at the center of the display in order to maintain a symmetric relation between the five lists. In essence, the list with the least important information for the purpose of the medication reconciliation is placed at the focus point of the user. Third, this design choice places the two similar-but-not-exact lists (2 and 3) on *opposite* sides of the exact-match list even though they represent the hardest to detect, discrepancies, which should be carefully compared and contrasted item by item. Finally, despite the prominent emphasis the authors give to the use of animation, they report that users had to view the animation multiple times in order to comprehend the spatial layout. One reason for the users' confusion might be that too many items are moving on the screen at the same time. Rather than helping users in the reconciliation process, the animation may actually distract them and shift their attention away from it.

IV. Design Objectives

To understand functional requirements and determine design objectives, we conducted group and individual interviews with Veterans Health Administration (VA), and Home Health (HH) nurses. We also observed providers as they reviewed HH plans of care. We noted that in complex cases involving high-risk discrepancies, the reconciliation process could spread over time and may include multiple phone calls and /or written communication between the physician's team and the HH nurse before a final reconciliation was achieved.

Providers stated that the reconciliation process is often tedious with potential for failure to detect or correct discrepancies. They expressed concerns regarding the validity of the data they receive from HH and the value of reconciling the medication lists on the CMS-485 paper form (since similar discrepancies would often reappear in subsequent communication with HH). This finding suggests that referring providers and HH nurses have different mental models of the *meaning* of each medication list and this consideration was brought into our design. Provider's main requests were for an electronic system, integrated with the Electronic Health Record (EHR) that would ensure accurate information, ease the comparison of medication lists, and allow them to take actions on each discrepancy while automatically updating the Electronic Health Record with the results of the reconciliation.

V. Design and Visual Concepts

Based on the functional requirements and design objectives

developed through the interviews and observations we constructed paper prototypes of the user interface. We continued to rely on principles of user-centered design [17][18] as we shared the paper prototypes with a small sample of VA physicians to iteratively refine and expand the design. The iterative refinement of the visual design and functionality continued throughout the development of the software prototype by ad-hoc testing with some of the same physicians who had interacted with the paper prototypes as well as physicians who were naïve to the initial design. Here we detail key design concepts and decisions, and elaborate on the reasons behind these decisions.

A. *Separation of concerns*

The role of a visual interface is to facilitate the user comprehension, decision making and action [19]. In this work, we focused on the design of medication reconciliation tool that use medication lists that have already been matched using an independent external pre-process. The separation between the matching algorithm and the visualization ensures that our system can be used in other settings that may impose other matching requirements. For the purpose of testing our software prototype (Section VII), we pre-processed the medications lists using a system developed by our collaborators at MITRE[20].

B. *First do no harm.*

Currently, when physicians work with the paper home health plans of care they can scan the list of medications very quickly. At times, physicians ignore medication discrepancies that they deem insignificant (e.g., over the counter medication or nutritional supplements) and may not update the VA records or order to discontinue use.

Although the paper-based reconciliation process is often inaccurate at time, it is nevertheless fast when there are few discrepancies or when none of the discrepancies seem to pose a significant risk. Based on this analysis we developed a *first do not harm* approach throughout the design process to ensure that handling the common case remains simple and fast with minimal requirements from the user. Complex situations are infrequent and by definition should require close attention and careful assessment by the physician. It is thus acceptable that handling complex situations will take longer time and require more complex interactions.

C. *"Above all else show the data"*

In his book "The Visual Display of Quantitative Data" [21], Edward Tufte introduces the concepts of Data-Ink and Data-Ink Ratio saying "*A large share of ink on a graphic should represent data-information*". Tufte codify them as the Laws of Data-Ink: Above all else show the data; Maximize the data-ink ratio; Erase non-data-ink; Erase redundant data-ink; and Revise and edit.

We employ the Laws of Data-Ink throughout the design to develop clear and concise visualization by identifying and removing redundant information (Section V.D), condensing annotations (Section V.D.4), reducing extraneous information (Section V.D.4), and designing user interactions that enable users to directly interact with visual items without the need for buttons and labels (Section VI.C).

D. *My List vs. the Other List*

Medication reconciliation has been defined by CMS[22] as "*The process of identifying the most accurate list of all medications that the patient is taking […] by comparing the medical record to an external list of medications obtained from a patient, hospital, or other provider"* (our emphasize)**.** Note that from the perspective of the physician there is a clear distinction between the two medication lists based on the physician's ownership of one of the lists. This distinction is particularly prominent in our setting because physicians are often unclear regarding the source and credibility of data in the medication list reported by HH, as we describe in Section IV.

The main theme of our design is centered on ownership and trust distinctions between the two lists, which we codify by the *'my list vs. the other list*' concept. We designed the visual interface specifically for the owner of the first medication list and use that list as the anchor for the rest of the display. For the '*other list'* we only present its discrepancies with respect to the owner list as shown in Figure 2**.**

1) *Discrepancies*

We identify four types of potential discrepancies,

- *No discrepancy*: This case can be further divided into exact and equivalent matches.
- *Prescription discrepancy*: discrepancies such as name, form, dose and frequency.
- *Missing at home*: indicate the *omission* of medication in accordance with our notion that the owner's list (VA in this case) is the definitive list for the patient.
- *Unexpected*: medications that are not on the owner's list. we further distinguish between three subcases: *Discontinued* and *Expired* indicating medications previously on the actively VA medication list, and *Unknown* indicating a medication for which no previous VA information is available.

There are various ways of ordering the medications, such as alphabetically or based on conditions. Condition based

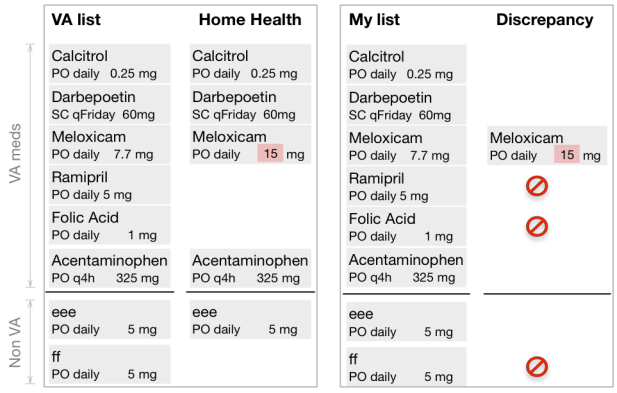

Figure 2 My List vs. the Other List: Two initial paper mockups, Left: the original two medication lists. Right: Depicting discrepancies relative to the physician list.

ordering can be beneficial in some setting although a medication can be associated with multiple conditions. Unfortunately, such information is not recorded in our setting. With no good contextual information to help in the ordering, we elected to order the medication based on the type of discrepancy and the display order in the current VA EHR (and alphabetically within each group). In particular, we use the following group ordering,

1. *Prescription discrepancy:* we presumed that these discrepancies (such as differences in dose and/or frequency) are the hardest to detect because they require close examination of differences between two text strings.
2. *Missing at home:* omission of a prescribed medication indicates a departure from the physician's medication regime.
3. *Discontinued*: may indicate erroneous use of previously prescribed medications or may only reflect the need to dispose of them.
4. *Expired*: may suggest the need to renew the prescription, the need for proper disposal, or erroneous use.
5. *Unknown* medications at home. These could represent medications prescribed by other physicians or over the counter medications.
6. *No discrepancy*: these medications are listed for completeness. When considering whether a given medication is appropriate the MD needs to consider the entire list of medications, not just those with discrepancies. They are listed last in our interface so that they user can focus on medications with discrepancies first.

2) Situational Awareness

During a medication reconciliation session, a physician needs to maintain a clear mental picture of the overall situation and the progress in the reconciliation process. This may include how many discrepancies exist, which ones have been addressed, whether they represent omission by a VA provider or represent medications prescribed by a non-VA physician (the actions for these two situations would be different).

The design outlined above supports the physician's situational awareness by clearly indicating where there are discrepancies and highlighting specific indicator mismatches. The vertical representation of the two lists enables the physician to quickly scan the display and maintain an updated and accurate mental picture of their progress in the reconciliation process. Additional visual cues to enhance situational awareness [23] are described in the following sections.

3) Handling unknown medications

Our initial approach was to separate medications found at home but not on the VA list ("unknown medications") from medications prescribed by a VA provider and treat them differently because they are not part of the owner's list (VA). A few physicians worried about polluting the VA list, which they considered the gold standard, with unimportant or insignificant supplements and over the counter medications. Upon further discussions with providers, we reached a consensus that in order to maintain consistency all the medications should be presented and dealt with in a single unified list.

4) Annotated list

In order to minimize information overload and maximize data to ink ratio, we initially indicated no discrepancy by leaving the entry in the discrepancy list empty. This decision was based on the principle that the discrepancy list should only be used to indicate that there is a discrepancy and describe it in simple terms. During the iterative design process, we found that physicians felt uncomfortable with the empty space concept and often wondered if that meant information was missing. In fact, some physicians expressed complete lack of trust in automated systems and wanted to see the exact wording of the prescription in order to confirm for themselves that there was no discrepancy, a demand that undermines the concept of providing cognitive support. To address these concerns, we introduced a checkmark to explicitly affirm no discrepancy was found.

Physicians were also uncomfortable with the use of icons to indicate that a medication was missing from the home. We thus annotated a missing medication entry with a label *'missing'* though we distinguish it from other medications by using italics and slightly faded the label to the background by using light gray color as shown in Figure 3.

It is possible that the owner of a medication list might have additional relevant information about items on that list, which can be used to further *annotate* the discrepancy list (Figure 3). Such information might include patient adherence, patient sensitivity, prescription status (expired or discontinued), or some previous communication with HH. One must be careful though not to introduce too much extraneous information that may clutter the display. We experimented with a simple paper clip icon to indicate that additional information is available but later opted to remove it for the sake of clarify and to reduce cognitive load. On the other hand, we found out that indicators such as 'stopped' and 'expired' were very helpful to physicians as they sought to understand the implication of finding a medication at home that is not on the VA list. Because these are annotations and not the medications themselves, we initially presented them to the right of the discrepancies. This caused the display to spread too wide forcing user to scan back and forth horizontally and thus reduce the effectiveness of the vertical scan (Section V.D.2). Our solution, shown in Figure 4, is to inline some of the annotations inside the owner (VA) list using non-obtrusive faded labels similar to the *'missing'* label. The inline annotations approach reinstates the vertical scan and emphasizes that these annotations represent the owner's (VA) own knowledge rather than an external source such as a verbal description by the patient (note that the clip was not in lined).

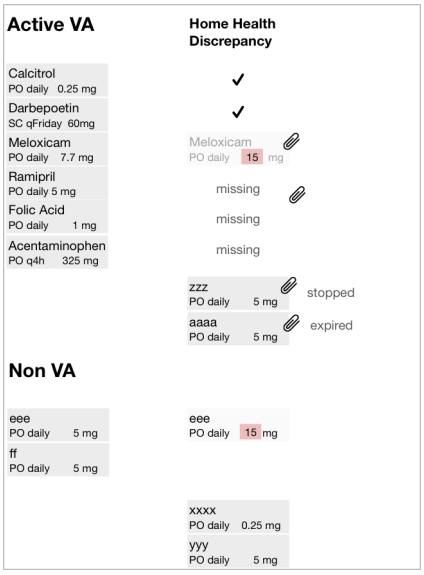

Figure 3 Annotating discrepancies: Using checkmarks to explicitly indicate no discrepancy and a 'missing' label to indicate medication not found at home. Other annotations such as 'stopped', and 'available notes' (paper clip) are placed to the right of the discrepancy.

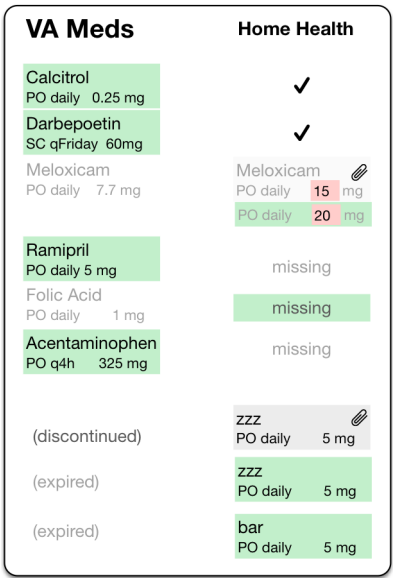

Figure 4 Final paper mockup. Using inline annotations to provide contextual information and preserve fast vertical scanning. Green background represent user decisions as described in the

E. On the use of animation

The physician's main goal is to address any existing discrepancies and it's the system goal to effectively present discrepancies to the physician. TwinList presents the user with the original two medication lists and then uses animation to show how the two lists are rearranged into five different lists based on the types of discrepancies. Our approach is to forgo this initial phase and immediately present the user with a concise and actionable representation of the discrepancies that need to be addressed. While the use of animations can be informative, animations impose time delays which are incongruent with our approach of making the average case fast.

VI. Descriptive vs. Declarative Orders

The design outlined above focuses on providing a clear and concise presentation of the discrepancies in order to improve situational awareness. An effective system must also facilitate the analytical reasoning process by supporting the analytical discourse between users and their information. In the following, we analyze the set of potential orders a physician may issue and show how we arrange them into four categories. We describe both declarative and descriptive based approaches for designing efficient and consistent interactions with the system.

A. Contextual orders

Consider a case where a medication was found at the patient home and the physician does not see a reason to stop or change its use. If this is a new medication than one potential order might be "Add to VA list". In a case of a known medication but with a different indication, the order might be "Update VA record". Similarly, if this is an expired VA prescription then a relevant order might be *"Renew prescription."* There are clear differences between these orders and each may require different actions yet they have a similar semantic, namely "Accept the Home Health finding and update the VA medication list". Moreover, these orders are mutually exclusive with regard to a particular discrepancy such that only one is viable.

Presenting a list of all possible orders, each with its own unique wording, will result in a cluttered and cumbersome interface. Even if we mark as disabled the orders that are not suitable for each discrepancy, it will force the physician to linearly scan each and every potential order and comprehend the differences between the orders before making a decision. Rather than empowering the physician, the system would only impose additional cognitive load and would slow the decision making process.

After reviewing all potential orders and their semantics, we identified four semantic groups such that only one order from each group is viable for each discrepancy. The four groups and their semantics are,

Enforce VA: Inform Home Health to request the patient adhere to the original VA order. This may imply changing or replacing the medication at the patient home or ensuring the patient does not take the medication.

Accept HH: The current medication (or lack thereof) at the patient home is acceptable. The VA records should be updated accordingly (i.e., add to, delete from, or edit dose or schedule in the VA medication list).

Clarify: This is a special case that may arise in a Home Health setting. In some cases, the discrepancies may require additional clarification. For example, if the VA's medication list notes that an outside physician has prescribed the patients a medication at a given dose and frequency (e.g. simvastatin 40mg once a day) but the home health list has different information (e.g. simvastatin 20mg once a day) the physician will likely want to ask the reason for this discrepancy such as 'does the patient take it differently than prescribed?' or "did the prescribing MD change the prescription?". Typically, the

physician will add a message to the home health nurse, asking for additional information or clarification.

No Action: The physician deems the discrepancy inconsequential and does not want to update the current order in the VA records.

B. Descriptive Interface

Traditional interfaces are based on descriptive actions which explicitly describe what should be done such as '*renew medication'* or '*update VA records'*. We use the grouping we define above to create a simple interface in each row consisting of four radio buttons arranged in a two by two box as shown in Figure 5. To maintain consistency, the two radio buttons on the left refer to the two most often used groups *Enforce VA* (top) and *Accept HH* (bottom). The radio buttons on the right represent the less frequent scenarios *Clarify* and *No Action.* We chose to place the *Enforce VA* and *Accept HH* vertically in order to minimize mouse movements and fit with the overall flow of scanning the medications vertically from top to bottom. Although the radio buttons arrangement is fixed in all the rows, the labels do reflect the specific appropriate order for the discrepancy in each row. For example, if the HH list contained a discontinued med the labels change to *'Do not use'*, and *'Renew'*.

A reconciliation process is not necessarily a linear process and the physician may temporarily skip some discrepancies in order to gain a better holistic view of all the discrepancies. We add a background color to each choice the physician has made in order to help the physician maintain an overall notion of which discrepancies he/she has already addressed and which discrepancies still stand during the reconciliation process. A selected choice for *'Enforce VA'* and *'Accept HH'* is marked by a green background, a *'Clarify'* is marked with yellow and *'No Action'* in blue.

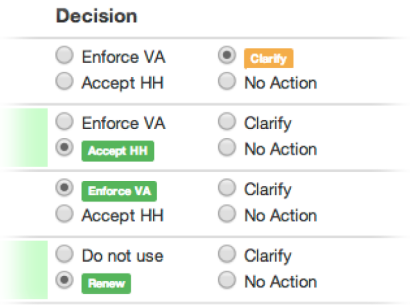

Figure 5 In a Descriptive Interface the physician explicitly states the action that should be taken. (portion of prototype screenshot).

C. Declarative Interface

In contrast to the descriptive actions, a declarative action describes what the results should be and leaves the exact details unspecified, such as 'the VA medication is the correct one'. In this context, the four groups we identified above form a set of four declarative orders, *VA*, *HH*, *No Action* and *Clarify*. We can map the first three to direct interactions with the display without the need for an explicit interface such as the radio buttons. Specifically, the physician selects (clicks on) an item in the VA (or HH) column to declare this the desired outcome. A click on a selected item unselects it and removes the order resulting in *No Action*. The physician can also directly switch from a VA order to a HH order and vice versa without the need to go through an intermediate *No Action* by selecting the desired radio button directly. We note that these declarative orders are unambiguous and translate to exactly one descriptive action based on the particular discrepancy. A selected item is marked with a green background (Figure 6).

A *Clarify* order is a statement that neither option is appropriate in the current context. It is a very special case because it also means the physician cannot sign the CMS485 form until the issue is resolved. Because of this unique case, we opted to use a dedicated button to mark a discrepancy as *Clarify* and mark both entries with a yellow background.

The declarative approach facilitates much faster and simpler interactions. The physicians view and interact directly with the medication and the discrepancy text rather than having to shift the attention to the radio buttons on the right, find and select the appropriate one and then shift attention back to the lists.

Though we believe this is faster and simpler interaction, we acknowledge that it is a departure from traditional interfaces. For this reason, we implemented both approaches in our prototype to look at their usage during the user study. We note that the TwinList interface can be regarded as a declarative interface although the semantic is somewhat different and the authors do not present it as such and do not discuss the implications.

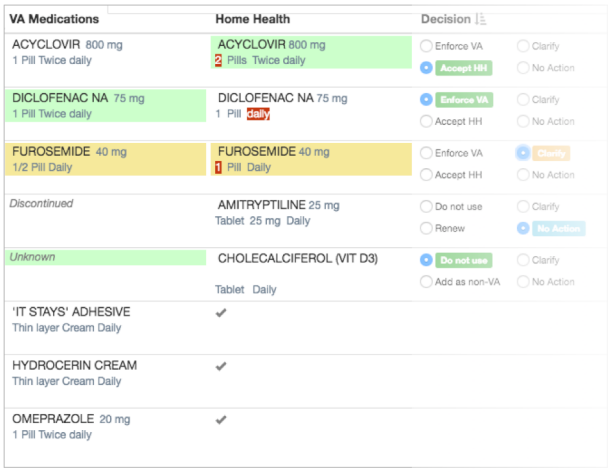

Figure 6 Declarative Interface (left two columns). A physician indicates the desired outcome by selecting (clicking) the VA or HH medication or discrepancy. In general, the *clarify* action, which is unique for the HH setting, requires an additional *clarify* button. (prototype screenshot)

VII. Visualization Prototype

We implemented our design as a web based medication reconciliation software prototype shown in Figure 7. The software prototype addresses additional requirements beyond the key design issues described above. In particular, the

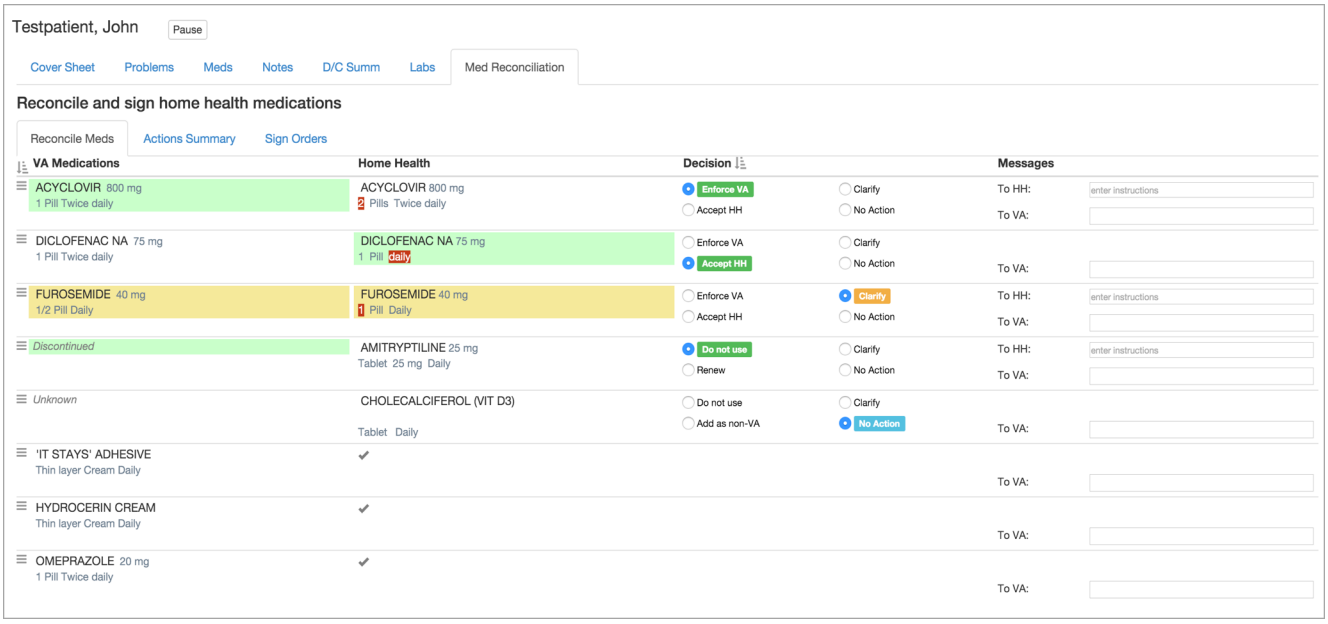

Figure 7 Screenshot of the final visual interface. In order to provide a contextual medical history for the user study, we developed an additional minimal simulated Electronic Health Records (EHR) that mimic the Computerized Patient Record System (CPRS) VA providers use regularly in their clinical practice (accessed via the tabs at the top).

interface enables physicians to add notes with additional instructions or questions to home health. The prototype also enables users to sort the medications alphabetically, based on the group each orders belongs to, or manually rearrange the rows by dragging them up or down the list. Additional screens provide the physician with a concise summary of the orders and a final signature screen, Figure 8, that depict the final medication list and orders the physician can sign.

VIII. User Study

We evaluated the effectiveness of our system and reported on that study and findings in another paper [11]; In the following, we summarize the key findings for completeness. To conduct the user study, we extended the software prototype and incorporated (limited) additional information screens that simulated information physicians have access to on a regular VA system. We employed repeated measures ANOVA to test the hypotheses that the system: 1) Improves *accuracy* by reducing the number of unaddressed medication discrepancies, 2) Improves *efficiency* by reducing the reconciliation time, 3) has good perceived *usability*.

Nineteen physicians with experience in managing home health referrals were recruited to participate in this within-subjects experiment. Participants completed medication reconciliation for one fixed 'warm up' and two randomly chosen clinical cases in each of two conditions. The first condition (paper-based) simulated current practice – reconciling medication discrepancies between a paper plan of care (CMS-485) and a simulated Electronic Health Record (EHR). For the second condition (electronic) participants used our medication reconciliation module, which we integrated into the simulated EHR.

The results support the improved accuracy hypothesis. Participants left more discrepancies unaddressed in the paper-based condition than the electronic condition, $p < 0.0001$ (Paper *Mean* = 1.55, *SD* = 1.20; Electronic *Mean* = 0.45, *SD* = 0.65). Our hypothesis that users would perceive good usability compared to the traditional CMS486 paper form is supported. Participants reported improved ability to detect and correct discrepancies with increased confidence in the results. The standard SUS score was 86.5 which correlates to an "excellent" rating [24][25]. The electronic process was the preferred process and was overwhelmingly well received.

Contrary to our efficiency hypothesis, participants took the same amount of time to complete cases in the two conditions. Based on participant comments and our post-hoc analysis of user's exploration of the simulated record, we hypothesized that an unintended (but highly beneficial) consequence of the electronic system was that participants were spending more

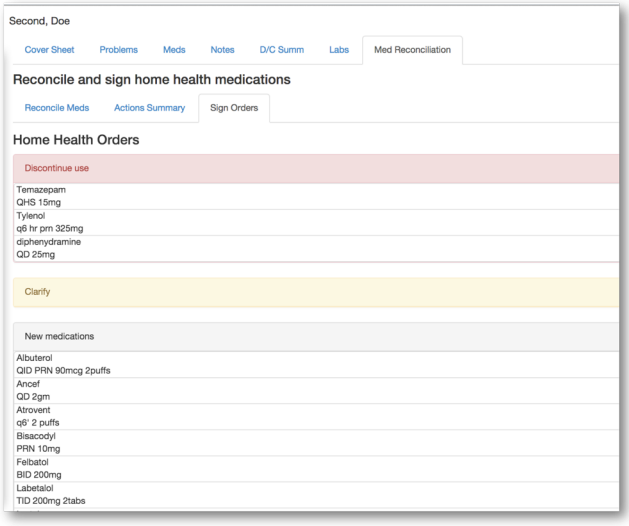

Figure 8 A sign orders screen depicting a concise summary of the doctor orders.

time foraging in the simulated EHR for information related to the *appropriateness* of medications. This post-hoc hypothesis is supported by examining the number of times participants switched between display panels (i.e., tab switches) in the simulated EHR during the paper-based vs. electronic scenarios. This need for a simple and direct access to the patient record during medication reconciliation is consistent with previous studies that noted providers' desire for interoperability between medication reconciliation tools and the patient records [26].

IX. Conclusion

In this paper, we present the design of an interactive visual interface for medication reconciliation. Our interface is intended to provide cognitive support for VA physicians treating patients receiving home health services. The two mechanisms by which our tool supports users is by making differences between "My list vs. their list" easy to detect, and by providing the capacity for the user to easily act on the discrepancies through the same interface. We described the iterative design process, the software prototype implementation and presented insights from our user study.

The main contributions of this work are in the discrepancy-first approach, the detailed discussions of key design decisions, and the principles behind them. We believe that while the home health setting introduces some unique challenges, the overall design of our medication reconciliation tool is directly applicable in other settings.

Limitation: the work presented here is based on the assumption that the two medication lists can be matched pairwise. For example, a prescription for a 20 mg twice a day may have been filled as 40 mg with instructions to break the pill. Additional complications can arise, such as if two medications in list A should be matched with two or three medications on list B. Although we have developed initial concepts for this complex edge case, we did not address these complex situations in this work.

References

[1] S. Kripalani, C. N. Theobald, B. Anctil, and E. E. Vasilevskis, "Reducing Hospital Readmission: Current Strategies and Future Directions," *Annu. Rev. Med.*, vol. 65, pp. 471–485, 2014.

[2] S. F. Jencks, M. V. Williams, and E. A. Coleman, "Rehospitalizations among Patients in the Medicare Fee-for-Service Program," *N. Engl. J. Med.*, vol. 360, no. 14, pp. 1418–1428, Apr. 2009.

[3] V. Mor, O. Intrator, Z. Feng, and D. C. Grabowski, "The revolving door of rehospitalization from skilled nursing facilities," *Health Aff. (Millwood)*, vol. 29, no. 1, pp. 57–64, 2010.

[4] D. M. Stitt, D. P. Elliott, and S. N. Thompson, "Medication discrepancies identified at time of hospital discharge in a geriatric population," *Am. J. Geriatr. Pharmacother.*, vol. 9, no. 4, pp. 234–240, 2011.

[5] L. Jeffs, M. P. Law, S. Straus, R. Cardoso, R. F. Lyons, and C. Bell, "Defining quality outcomes for complex-care patients transitioning across the continuum using a structured panel process," *BMJ Qual. Saf.*, vol. 22, no. 12, pp. 1014–1024, Dec. 2013.

[6] K. T. Unroe, T. Pfeiffenberger, S. Riegelhaupt, J. Jastrzembski, Y. Lokhnygina, and C. Colón-Emeric, "Inpatient medication reconciliation at admission and discharge: A retrospective cohort study of age and other risk factors for medication discrepancies," *Am. J. Geriatr. Pharmacother.*, vol. 8, no. 2, pp. 115–126, Apr. 2010.

[7] P. Varkey, J. Cunningham, J. O'Meara, R. Bonacci, N. Desai, and R. Sheeler, "Multidisciplinary approach to inpatient medication reconciliation in an academic setting," *Am. J. Health-Syst. Pharm. AJHP Off. J. Am. Soc. Health-Syst. Pharm.*, vol. 64, no. 8, pp. 850–854, Apr. 2007.

[8] "Medication Reconciliation During Transitions of Care as a Patient Safety Strategy: A Systematic Review | Annals of Internal Medicine." [Online]. Available: http://annals.org/article.aspx?articleid=1656444. [Accessed: 04-Sep-2016].

[9] A. Brody *et al.*, "High Prevalence of Medication Discrepancies Between Home Health Referrals and Centers for Medicare and Medicaid Services Home Health Certification and Plan of Care and Their Potential to Affect Safety of Vulnerable Elderly Adults," *J. Am. Geriatr. Soc.*

[10] C. F. Corbett, Daratha K. B., and et all, "Nurse identified hospital to home medication discrepancies: Implications for improving transitional care," *Geriatr. Nurs.*, vol. 31, pp. 188–196, 2010.

[11] H. S. Kramer, B. Gibson, Y. Livnat, I. Thraen, A. Brody, and R. Rupper, "Evaluation of an Electronic Module for Reconciling Medications in Home Health Plans of Care," *Appl. Clin. Inform.*, vol. 7, no. 2, pp. 412–424, 2016.

[12] Joint Commission on Accreditation of Healthcare Organizations, USA, "Using medication reconciliation to prevent errors," *Sentin. Event Alert Jt. Comm. Accreditation Healthc. Organ.*, no. 35, pp. 1–4, Jan. 2006.

[13] "Hospital-Based Medication Reconciliation Practices: A Systematic Review | Jul 23, 2012 | JAMA Internal Medicine | JAMA Network." [Online]. Available: http://archinte.jamanetwork.com/article.aspx?articleid=1203516. [Accessed: 04-Sep-2016].

[14] E. G. Poon *et al.*, "Design and implementation of an application and associated services to support interdisciplinary medication reconciliation efforts at an integrated healthcare delivery network," *J. Am. Med. Inform. Assoc. JAMIA*, vol. 13, no. 6, pp. 581–592, Dec. 2006.

[15] C. Plaisant *et al.*, "Twinlist: Novel User Interface Designs for Medication Reconciliation," *AMIA. Annu. Symp. Proc.*, vol. 2013, pp. 1150–1159, Nov. 2013.

[16] C. Plaisant, J. Wu, A. Z. Hettinger, S. Powsner, and B. Shneiderman, "Novel user interface design for medication reconciliation: an evaluation of Twinlist," *J. Am. Med. Inform. Assoc. JAMIA*, vol. 22, no. 2, pp. 340–349, Mar. 2015.

[17] K. Holtzblatt, J. B. Wendell, and S. Wood, *Rapid contextual design:A how-to guide to key techniques for user-centered design*. San Francisco, CA: Morgan Kaufmann, 2005.

[18] J. Nielsen, *Usability engineering*. Elsevier, 1994.

[19] J. J. Thomas and K. A. Cook, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE CS Press, 2005.

[20] D. Tresner-kirsch *et al.*, "Automated Detection of Discrepancies in Pharmaceutical Order-Entry Data.," *Rev.*

[21] E. Tufte, *The Visual Display of Quantitative Information*, 2nd Edition. Graphics P, 2001.

[22] "Medication Reconcilliation." Centers for Medicare & Medicaid Services (CMS).

[23] M. R. Endsley, *Designing for situation awareness: An approach to user-centered design*. CRC, 2009.

[24] J. Brooke, "SUS-A quick and dirty usability scale," *Usability Eval. Ind.*, vol. 189, no. 194, pp. 4–7, 1996.

[25] A. Bangor, P. Kortum, and J. Miller, "Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale," *J Usability Stud.*, vol. 4, no. 3, pp. 114–123, May 2009.

[26] H. Monkman, E. M. Borycki, A. W. Kushniruk, and M.-H. KUOa, "Exploring the Contextual and Human Factors of Electronic Medication Reconciliation Research: A Scoping," *Context Sensitive Health Inform. Hum. Sociotechnical Approaches*, vol. 194, p. 166, 2013.

Network analysis of treatment patterns in breast cancer care

Shane Weisberg, Douglas Hill, Rebecca Faill, and Amar K. Das

***Abstract*—Though cancer treatment is standardized at the national level, not much of what we know about treatment patterns of cancer patients is well supported by data. Cancer patients are diagnosed and undergo complicated treatments over extended periods of time and see new providers with different specialties at each step along the way. To find information about general patient care patterns, we constructed a patient-sharing network of breast cancer providers and patients using provider-patient encounter data and observed how this network changes over the duration of treatment of those patients. By using visualizations in three-dimensions in unison with mathematical descriptions of these networks, we are able to recognize patterns in how patients enter and progress through the network of providers.**

***Index Terms*—breast cancer, patterns of care, network visualization, centrality**

I. Introduction

With the introduction of big data into medical care, tools are now available to answer some of healthcare's most fundamental questions, such as whether there are recognizable patterns of care in medical practices and, should they exist, whether they match common conceptions of what they should be. This paper will use electronic health records from cancer care centers in New Hampshire and Vermont to construct a dynamic network of providers based on patient-sharing relationships between providers and use both three-dimensional visualizations and traditional network statistics to detect and identify patterns in cancer care.

For this study, we are interested in how this network of providers changes over time as patients enter and exit the network and whether patterns of care arise in this network. We find that a provider's location serves as the dominant structure within the network with respect to shared patients, which is to say that providers practicing at the same location share more patients. Beyond this macro structure, we also find that a provider's specialty can serve as a predictor of when in the cancer care cycle a provider will be most central with respect to the treatment of patients. Specifically, we will look at eigenvalue centrality to see which types of providers are most important to the network [1]. Those providers which are highly connected are important in that they see many patients and can serve as a focal point for resources but also act as a bottleneck in patient care as many physicians and surgeons often restrict the number of patients they can see in a day [2]. Identifying these central providers is important to making patient care more efficient and sustainable.

The data used for analysis herein is provider-patient encounter information gathered over two years from the Dartmouth Hitchcock Medical Center (DHMC) in Lebanon, New Hampshire and neighboring institutions in and around Manchester, New Hampshire and Nashua, New Hampshire. Specifically, we will focus on female breast cancer patients and their paths as they enter the vast network of providers and progress through their treatments. Cancer care is ideal for this type of study as it involves visits with many different specialists at different times throughout the treatment process and should be fairly uniform across different treatment centers.

In the following sections, we will first take a look at past work related to this topic before we discuss the data used for this study as well as the computational tools and methods used to analyze that data. Subsequently, we will discuss our results, and display sample visualizations of the networks and corresponding network statistics that solidify the results. We will finish by discussing our conclusions and examining possible extensions where this methodology could be used.

II. Previous Work

The software used for visualizing our network was built for this purpose, and though similar networks of providers and patients have been studied, our methodology for doing so is new, especially with the addition of a time dimension to the network analysis. Bridewell and Das [3] used a similar network of providers at the Palo Alto Medical Foundation and Stanford Hospital to determine that while the institutions are closely linked, patients tend to stay at the site where they initially received care. The Bridewell and Das study viewed the network as a static object rather than studying it over time.

Katz et al. [4] did similar analysis of patterns of patient movement looking at referrals for breast cancer patients. In this study, however, regression models rather than a network model were used to support their findings. Wang et al. [5] used network analysis and specifically different measures of network centrality to study cancer treatment, though their focus was gene signatures rather than patient care patterns. Unnikrishnan et al. [6] looked at relationships over time between providers in intensive care units, but used radically

Shane Weisberg is with the Department of Biomedical Data Science, Geisel School of Medicine, Hanover, NH, USA, e-mail: shane.c.g.weisberg.16@dartmouth.edu.

Douglas Hill is with the Department of Biomedical Data Science, Geisel School of Medicine, Hanover, NH, USA, e-mail: douglas.p.hill@dartmouth.edu.

Rebecca Faill is with the Department of Biomedical Data Science, Geisel School of Medicine, Hanover, NH, USA, e-mail: rebecca.faill@dartmouth.edu.

Amar Das is the director of Biomedical Informatics, Geisel School of Medicine, Hanover, NH, USA, e-mail: amar.dasl@dartmouth.edu.

different methods to support their claims. Keating et al. [7] also used networks to study relationships between providers, in this case all at the same hospital, however this study deals with transfers of information and ideas on a small scale rather than referrals of patients at the scale used in this analysis.

III. Data

The data used for this study was drawn from a patient cohort consisting of all female patients over 21 years of age who had a primary diagnosis of stage I-IV breast cancer as reported by the Dartmouth Hitchcock tumor registry. Diagnosis dates of patients in the sample range from April 2, 2011 to November 1, 2013 and care and treatment are recorded through year end 2014. Our sample contains 723 providers and 925 patients. Patients appear in anywhere from one to 421 rows in our dataset, where each billable piece of an encounter between a patient and a provider gets its own row. The median number of encounters of patients in our dataset is 19.

Each row of data details a specific interaction between a provider and a patient and gives information about the provider such as his/her location, specialty and the treatment given. For each billing code or provider-patient encounter we have the date of the interaction, the ID of the provider, the name of the department, the provider's specialty, the ID of the patient as well as a short description of the event. The provider and patient identification columns were used to determine which providers were connected in the network. The dates of the interactions for each patient were jittered uniformly by a random number of days less than or equal to 45. From this point, we normalized the dates such that each patient enters the network simultaneously. We identified the first appearance of each patient and subtracted this day number from each of the patient's appearances to standardize the time span over all patients.

Beyond manipulating dates, we also generalized both a provider's location and his/her specialty into broader categories. Provider locations were sorted into one of four categories—Upper Valley (Lebanon, Hanover, Plymouth and Lyme in New Hampshire and St. Johnsbury in Vermont), Manchester (and Bedford and Derry, New Hampshire), Nashua (and Hudson, Milford and Merrimack, New Hampshire) and Other. Some providers practice in multiple locations; we chose the most common location for our categorizations. Additionally, each provider was sorted into one of nine groups based on one of 50 unique specialties seen in the dataset. These nine groups with group size in parentheses are medical oncologists (71), medicine specialists (239), mental health specialists (10), ob/gyn specialists (46), primary care specialists (159), radiation oncologists (10), radiology specialists (17), surgical oncologists (8) and surgical specialists (163).

IV. Methodology

Using the data as outlined above, we built a network where providers are nodes and are linked by an edge if they share one or more patients in a given time span. We chose to represent our network using adjacency matrices, and Python was used to convert the rows of data into this form. Representing the networks as such makes both visualizing the network and performing calculations of network statistics simple. Given an adjacency matrix A and providers $p \in \{p_1, p_2, ..., p_{723}\}$ we say providers p_i and p_j are adjacent if they share one or more patients in a given time and the entry A_{ij} contains the number of patients shared by providers p_i and p_j. It is important to note that these relations are non-directional, and so this matrix is symmetric.

In order to observe changes in this network over time, it was necessary to expand this familiar structure to incorporate the dimension of time. Thus, we built a three-dimensional structure constructed of layers of adjacency matrices where each individual layer can be interpreted as the matrix A described above, and each layer represents the state of the network at a given time. So now, entry A_{ijt} contains the number of patients shared by providers p_i and p_j at time t.

This new structure makes defining a relationship non-trivial. In a static network, any patient who sees two providers at any point in their treatment becomes an edge between those providers in the network. But adding time complicates this relationship. We decided that a patient shall link two providers if and only if she sees both providers within a period of 30 days. We determine a patient to be seeing two providers at the same time if she sees both within a month. When building a layer of the network, we first filter the data to select only data in a 30 day interval and construct the adjacency matrix from that subset of the data. If a patient saw provider p_i on day d and provider p_j on day $d + t$ where $0 \leq t < 30$ then and edge connecting the two providers was added to the network on layers $\{d + t - 30, d + t - 29, ..., d, ..., d + t, ..., d + 30\}$ where each day in the simulation is given a layer. This large structure was then converted to the proper xgmml format to be read by our 3D visualizer and additionally printed in matrix form to perform matrix calculations in MatLab.

The 3D visualizer was built using the Unity Game Engine [8] and enabled us to build a dynamic, interactive network from our data. This network was built using a force-directed algorithm where nodes repulse each other and edges pull nodes towards each other with a force proportional to the weight of the edges. This structure is set up so that highly interconnected groups or clusters are displayed together in the graph and highly connected nodes tend to fall towards the center of the graph.

V. Results

Using visualizations and network statistics, we found that there are recognizable patterns in cancer care and that largely these patterns agree with clinical expectations of care patterns. This section will include screenshots of our provider-patient network generated with our 3D visualizer and will back up claims of structure using descriptive statistics of the network.

In Figure 1, the entire unfiltered network on the day all patients enter the simulation is displayed. Nodes are colored by a provider's primary location: maroon represents the Upper Valley; blue represents Manchester; gray represents Nashua. The image makes apparent that providers based in the same location share more patients and therefore are more closely

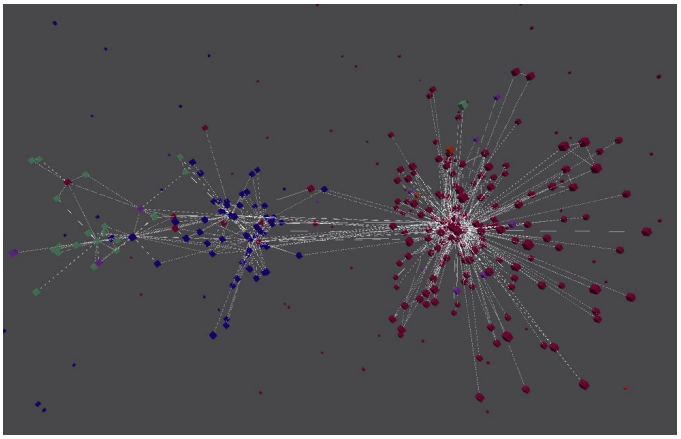

Figure 1. A snapshot from the 3D visualizer showing the unfiltered network of providers and patients on the first day of the time span. Nodes are colored by the primary location of the provider. Maroon nodes are providers who practice in the Upper Valley, blue nodes are providers who practice in Manchester and gray nodes are providers who practice in Nashua.

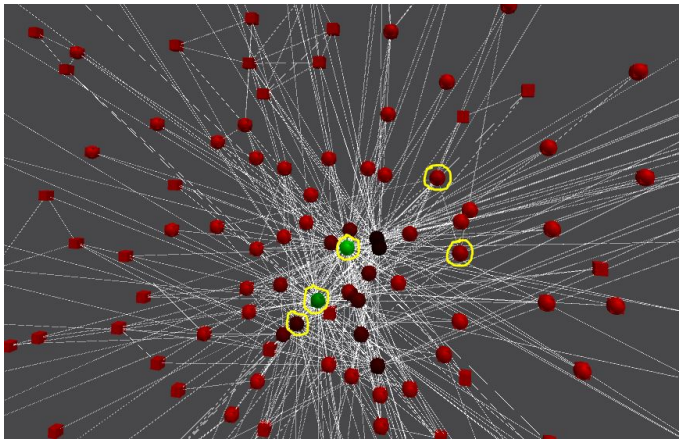

Figure 2. A snapshot from the 3D visualizer showing a two-dimensional projection of the unfiltered network on the first day of the time span. Nodes are colored on a scale from red to green where the reddest nodes have the fewest edges and the greenest nodes have the most edges. Nodes circled in yellow are surgical oncologists. Spherical nodes are those adjacent to the center circled node.

related. We can use density, a measure of concentration of edges to support this conclusion [9]. Density measures the ratio of the number of edges in a graph to the total number of edges possible in a graph. For an undirected network with n nodes, there are $\frac{n(n-1)}{2}$ possible edges. Table 1 shows number of edges, potential edges and densities for the full network pictured below as well as for networks that have been filtered by location. As expected, each of the filtered networks has a higher density than the full network, which further demonstrates the clustering of the network by location. The increase in density for the Upper Valley network is small since it is much larger than the Manchester and Nashua networks, and so adding these nodes and edges does not dramatically change properties of the network. This clustering in itself is not a groundbreaking discovery, however it does support the conclusions given by Bridewell and Das [3], and it would be neglectful to discuss this network without mentioning its primary structure.

Table I
DENSITIES OF NETWORKS FILTERED BY PRIMARY LOCATION

Location	All	Upper Valley	Manchester	Nashua
Number of edges	771	529	98	41
Possible edges	36,046	18,145	1378	325
Density	0.0214	0.0292	0.0711	0.1262

From here we will examine provider specialties and how the importance or centrality of different types of providers varies with time. Figure 2 shows a two-dimensional projection of the center of the full network on the first day range of the simulation where nodes are now colored on a scale from red to green where redder nodes have fewer adjacencies. Observable is that the surgical oncologists, which are circled in yellow, include both of the most connected nodes and are adjacent to most of the other central nodes. Nodes adjacent to the middle surgical oncologist are represented as spheres rather than cubes. Most patients see a surgical oncologist soon after their diagnosis as illustrated in the figure below.

Much can be learned about patterns in cancer care by looking at these networks in this manner and at how the networks change over time. However, a more efficient way to determine network trends and structure, specifically centrality [1], is by using network statistics, which are more exact and can be calculated and plotted over time. Figure 3 shows the average centrality score for four groupings of providers over the first 180 time windows of patient care for all providers and patients in the sample. The averages are smoothed using locally weighted linear regressions. We look only at the first 180 days of the time span as beyond this visits become less frequent due to more variability in scheduling, so the graphs become noisy and the centrality averages have less meaning.

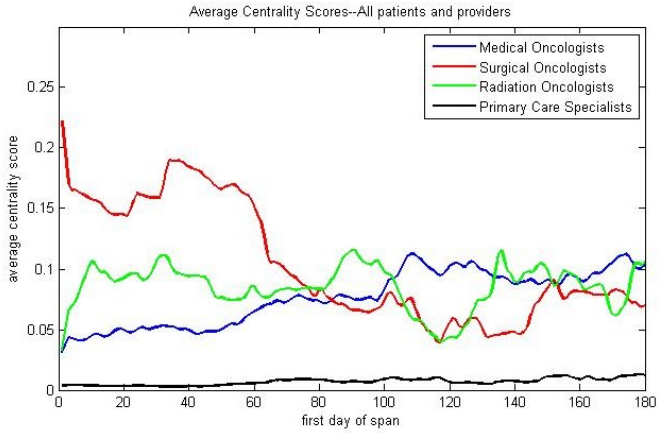

Figure 3. Average eigenvalue centrality over first 180 day ranges of the simulation for medical oncolologists (blue), surgical oncologists (red), radiation oncologists (green) and primary care specialists (black). Curves were smoothed using a locally weighted linear regression.

As we saw in the visualization, surgical oncologists are most central to the network on day one and they remain central over the first months of the time span. Surgical oncologists show a decreasing trend overall as radiation oncologists and

medical oncologists begin to fall towards the center of the graph as we approach the second half of the time span. Additionally, we see a very characteristic dip in the average centrality score of surgical oncologists in the first 40 days, which will be addressed later. Medical oncologists show a gradually increasing trend while the radiation oncologists have higher variability but trend neither upwards nor downwards. Primary care specialists are also shown in this plot. More than one fifth of the providers in the sample are primary care specialists. Therefore, no providers in this group are highly central to the network; rather they tend to float around the periphery of the network.

Additional insight, especially regarding the role of radiation oncologists, can be seen by filtering patients by the stage of cancer exhibited upon diagnosis4. In Figure 4, we return to visualizations of the network to observe this. In the snapshot below we show the network for the time span between days 21 and 51 filtered to include only stage I cancer patients. Again the nodes are colored from red to green with increasing number of edges, and this time the radiation oncologists are circled in yellow and the spherical nodes are all nodes connected to the topmost circled node. In this image, we see a pair of radiation oncologists right at the center of the network and these are the two most highly connected providers in the network, but we do not see a corresponding increase in centrality in Figure 3 as we would expect. A similar phenomenon can be seen in the network filtered by stage II patients in the day 10 to 40 span.

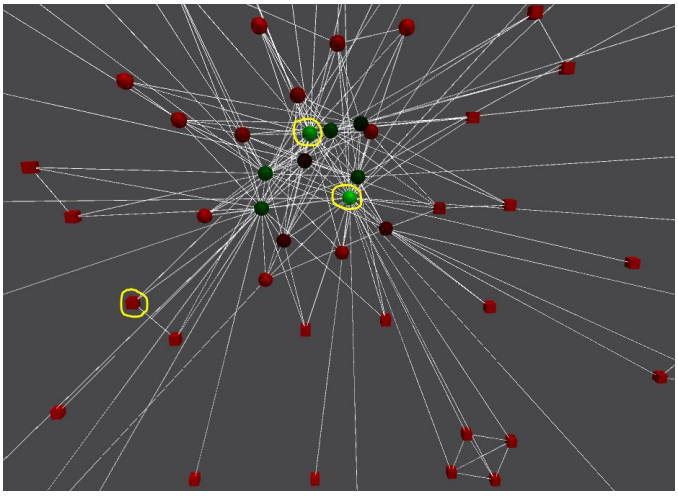

Figure 4. A snapshot from the 3D visualizer showing a two-dimensional projection of the network filtered to show only stage I cancer patients. Nodes are colored on a scale from red to green where the reddest nodes have the fewest edges and the greenest nodes have the most edges. Nodes circled in yellow are radiation oncologists. Sperical nodes are those adjacent to the uppermost circled node.

Again we can make centrality plots to resolve this issue and to better observe patterns of care over time. Figures 5 and 6 show the same information as Figure 3 after filtering the data by the patient's stage of cancer. There were only 89 stage III and 24 stage IV patients as compared 519 patients with stage I cancer and 274 patients with stage II cancer. Due the small sample size of the stage III and IV patients — and the highly noisy nature of the plots — they have been omitted here.

We see that surgical oncologists and medical oncologists show mostly the same trends as in Figure 2. However, the radiation oncologists show radically different behavior at the beginning of the time span as predicted by the visualizations. Noticeable is how in both plots, the green line representing radiation oncologists spikes above the red line for surgical oncologists in the first 20 days. Because this spike occurs at different times, it destructively interferes with itself and so is not shown in Figure 3. But we can see both in the filtered centrality plots and in the visualizations that radiation oncologists are highly central to the network, especially in the first month of care. We do not see this destructive interference pattern with other providers. Both graphs also show a third surgical spike around the day 140 to 170 span.

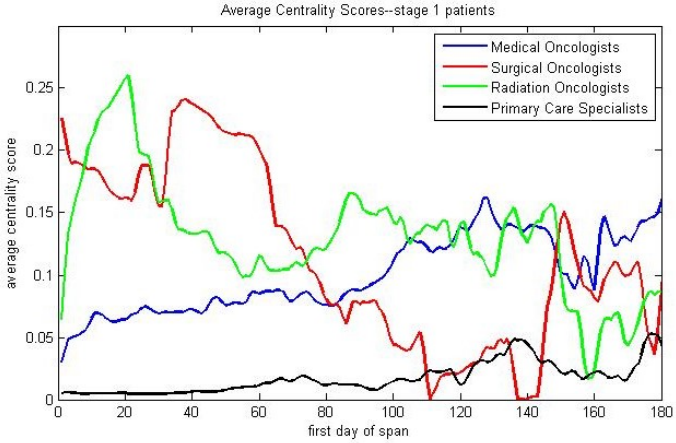

Figure 5. Average eigenvalue centrality over first 180 day ranges of the simulation for medical oncologists (blue), surgical oncologists (red), radiation oncologists (green) and primary care specialists (black) after data was filtered to include only stage I cancer patients. Curves were smoothed using a locally weighted linear regression.

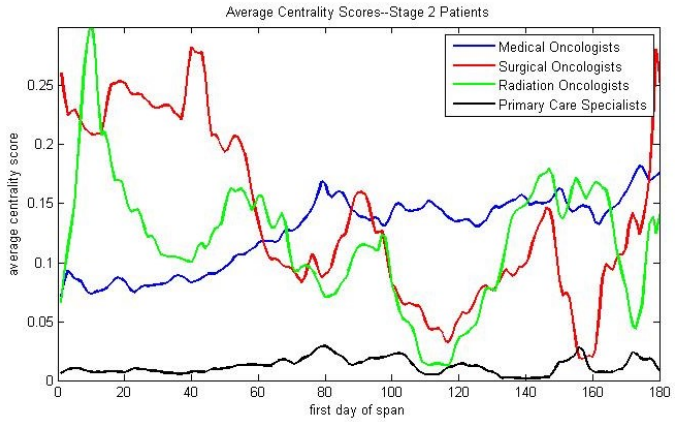

Figure 6. Average eigenvalue centrality over first 180 day ranges of the simulation for medical oncologists (blue), surgical oncologists (red), radiation oncologists (green) and primary care specialists (black) after data was filtered to include only stage II cancer patients. Curves were smoothed using a locally weighted linear regression.

Primary care specialists see a boost in their average centrality scores across the time span in the stage II plot when

compared to the stage I plot and the unfiltered plot. This boost most likely arises due to the shape of the centrality distributions exhibited by the primary care specialists. This group of providers has a within-group distribution of centralities that is highly skewed to the right, meaning that most of the providers have very low centrality scores while a few have very high scores. The distribution of centralities of primary care specialists has average skewness of 5.05—much higher compared to surgical and radiation oncologists, for example, which have skewnesses of 1.27 and 1.60 respectively. Skewness is calculated as $\frac{E(x-\mu)^3}{\sigma^3}$ where μ is the mean on the distribution of x and σ is the standard deviation of the distribution of x. This skewness accounts for the increase in centrality scores as with a decreased sample size, the highly central nodes have more influence on the average centrality. This suggests that while most primary care specialists are peripheral to the network, especially for stage I patients, a few select specialists become more central for care of stage II patients. With a larger sample of higher stage patients, it would be possible to asses whether this pattern continues for stage III and IV patients.

VI. Conclusions

There is a clear trend to be found in the network visualizations shown especially vividly in Figures 2 and 3. Patients enter the network on the first day and many see surgeons within the first 10 days after their diagnosis. After this initial surgical visit, patients may see a radiation oncologist, and the data suggests that there is more urgency to see a radiation oncologist for stage II patients than for stage I patients. This centrality spike for radiation oncologists likely comes from neoadjuvant radiation therapy, which is prescribed by surgeons to reduce or downstage a tumor before surgery [10]. Not surprisingly, we then see a corresponding spike in centrality of surgical oncologists as the actual surgery to remove the tumor is performed. From here, treatment is passed to radiation oncologists and medical oncologists who perform a combination of radiation therapy and chemotherapy. To draw further conclusions about treatment patterns past this early stage of cancer treatment, more data is necessary to ensure more coincidence between patients on later days in the time span.

This information about the early steps in breast cancer treatment can potentially have a variety of applications for hospitals, which can use this data to better allocate resources and time to make treatment more efficient. The centrality of the surgical oncologists and radiation oncologists can partly be attributed to the fact that these types of providers are scarce in the sample. Patients who need surgery or radiation must go through one of a very limited number of providers. Therefore, this method could potentially be used to diagnose resource shortages both at specific hospitals and in the health care system in general.

Further analysis can be performed on this data, including using different network statistics such as assortativity, from which we can learn a lot about referral patterns. Assortativity [11] is a measure of which types of nodes link to other types of nodes. A high assortativity rating demonstrates that nodes of high degree link to other nodes of high degree, which would suggest a network dominated by few providers while a low assortativity rating demonstrates that nodes of high degree link to nodes of lower degree, which would suggest a more uniform network. In the context of cancer care, a highly assortative network would imply that highly central providers refer patients to the same providers consistently rather than spreading out their referrals meaning that it is likely that the network is not being travelled by patients as efficiently as possible.

Similar analysis can be performed on datasets from other treatment centers. As cancer care is highly regularized nationally, we would expect that the same trends can be seen regardless of location. Our technique also has the potential to address more and broader questions about cancer care than the select and narrow one set forth here.

References

[1] Borgatti, Stephen P. and Martin G. Everett. "A Graph-theoretic perspective on centrality." Social Networks 28 2006, pp. 466-484.

[2] Gupta, Diwakar and Brian Denton. "Appointment Scheduling in Health Care: Challenges and Opportunities" 2008.

[3] Bridewell, Will, and Amar K. Das. "Social Network Analysis of Physician Interactions: The Effect of Institutional Boundaries on Breast Cancer Care." AMIA Annual Symposium Proceedings, 2011, pp. 152–160.

[4] Katz, SJ, Hofer TP, Hawley S, et al. Patterns and correlates of patient referral to surgeons for treatment of breast cancer. Journal of Clinical Oncology, 2007, pp. 271–276.

[5] Wang, Jianxin et al. "Integration of breast cancer gene signatures based on graph centrality." BMC Systems Biology 5, 2011.

[6] Unnikrishnan, K.P., Debprakash Patnaik, and Theodore J. Iwashyna. "Spatio-Temporal Structure of US Critical Care Transfer Network." AMIA Summits on Translational Science Proceedings, 2011, pp. 74–78.

[7] Keating, Nancy L. et al. "Factors Affecting Influential Discussions Among Physicians: A Social Network Analysis of a Primary Care Practice." Journal of General Internal Medicine, 2007, pp. 794–798.

[8] https://unity3d.com/

[9] Ester, Martin et al. "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," 1996.

[10] Urquhart, Robin et al. "How do surgeons decide to refer patients for adjuvant cancer treatment? Protocol for a qualitative study," 2012.

[11] Holme, Petter and Jing Zhao. "Exploring the assortativity-clustering space of a network's degree sequence" 2006.

Shane Weisberg is a student at Dartmouth College studying applied mathematics and computer science while working in the Department of Biomedical Data Science at the Geisel School of Medicine in Hanover, NH.

Douglas Hill is an associate staff member at the department of Biomedical Informatics at the Geisel School of Medicine in Hanover, NH with a focus in 3D data visualization.

Rebecca Faill is the Data Science Project Director at the Department of Biomedical Data Science at the Geisel School of Medicine in Hanover, NH.

Amar Das is an Associate Professor of Biomedical Data Science in the Geisel School of Medicine at Dartmouth College and is the Director of the Division of Biomedical Informatics in the Department of Biomedical Data Science.

m-TSNE: A Framework for Visualizing High-Dimensional Multivariate Time Series

Minh Nguyen[1], Sanjay Purushotham, PhD[1], Hien To[1], Cyrus Shahabi , PhD[1]
[1]University of Southern California, Los Angeles, CA, USA

Abstract

Multivariate time series (MTS) have become increasingly common in healthcare domains where human vital signs and laboratory results are collected for predictive diagnosis. Recently, there have been increasing efforts to visualize healthcare MTS data based on star charts or parallel coordinates. However, such techniques might not be ideal for visualizing a large MTS dataset, since it is difficult to obtain insights or interpretations due to the inherent high dimensionality of MTS. In this paper, we propose "m-TSNE": a simple and novel framework to visualize high-dimensional MTS data by projecting them into a low-dimensional (2-D or 3-D) space while capturing the underlying data properties. Our framework is easy to use and provides interpretable insights for healthcare professionals to understand MTS data. We evaluate our visualization framework on two real-world datasets and demonstrate that the results of our m-TSNE show patterns that are easy to understand while the other methods' visualization may have limitations in interpretability.

1 Introduction

Big data analytics in healthcare is emerging as a large amount of health and medical data are being generated every day. Recent development of different types of health sensors and e-health platforms has opened up great opportunities for collecting, monitoring and analyzing patients' health conditions from multiple data sources [1]. Performing analytics and extracting insights on healthcare data is challenging due to the large volume, high dimensionality, heterogeneity, and dynamic nature of the healthcare data. To address these challenges, many studies are being conducted in several fields such as machine learning, data mining, statistics, health informatics, etc. Data visualization is one such field that provides tools for visual interpretations of the underlying data patterns and trends, and helps in further data analysis.

Several techniques [2, 3, 4, 5] have been developed to visualize and analyze time series data. Since these techniques usually analyze univariate time series (UTS) data by handling only one data variable at a time (e.g., monitoring respiratory rate or heart rate over time to detect anomaly), they may not fully capture the inherent correlations of the multivariate time series (MTS) data [6]. Therefore, we believe that MTS data corresponding to multiple variables should be treated as a whole, rather than being broken into individual UTS as they can provide greater insight into data representing patients' conditions. However, multivariate data is notoriously hard to represent because of the difficulty of mentally picturing data in more than three dimensions [7].

Previous visualization efforts [8, 9, 10, 11] on multivariate data focused on displaying multiple dimensions of the data in 2-D plot, and left the interpretations to human observer. Figures 1 and 2 show the star chart and parallel coordinates visualizations that have been previously proposed for healthcare data visualization [10, 11]. These techniques might not be ideal for real-world applications which have high-dimensional MTS data points since they tend to result in complex visualization plots with less interpretability. Researchers have explored dimensionality reduction techniques [12, 13, 14] to handle high-dimensional data points by projecting them into low-dimensional space. These approaches have been applied successfully in computer vision applications where data consists of images. However, such techniques, in general, cannot be applied directly on medical MTS datasets because each data source may have different properties, e.g. time dependency structure is present in vital signs time series. Hence, there is a need for developing new techniques to visualize high-dimensional MTS data.

In this paper, we propose m-TSNE (Multivariate Time Series t-Distributed Stochastic Neighbor Embedding): a framework for visualizing MTS data in low-dimensional space that is capable of providing insights and interpretations of the high-dimensional MTS datasets. m-TSNE first calculates the similarity between

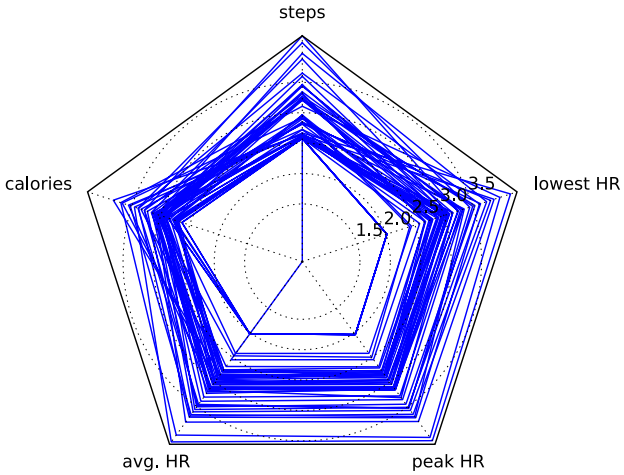

Figure 1: Star chart visualization [8] of MTS human monitoring data of a subject in ATOM-HP dataset (Section 4.1). Data has 5 variables: step counts, calories, lowest heart rate, average heart rate, and peak heart rate. Each polygon is a multivariate data point at a time instance.

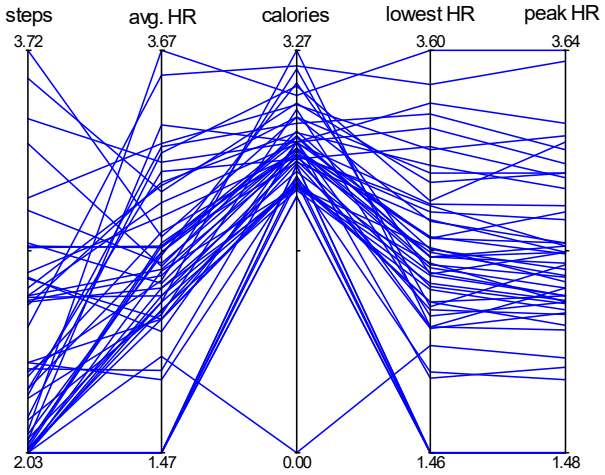

Figure 2: Parallel coordinates visualization [9] of the same data shown in Figure 1. Each polyline that connects the parallel axes represents a multivariate data point at a time instance. The number of polyline in the figure is the number of time instance in the MTS.

each MTS data points in high-dimensional space, based on Extended Frobenius norm (EROS) [6] which is a similarity metric for MTS data. Then, it computes the low-dimensional (2-D or 3-D) projection of the MTS data points using a gradient descent method by preserving the similarity relation between pairs of high-dimensional points. We conduct visualization experiments on two healthcare datasets: ATOM-HP dataset (Section 4.1), a dataset collected to study impact of chemotherapy on patient's activity; and Electroencephalogram (EEG) dataset [15], a dataset collected to study whether there is a genetic predisposition to alcoholism. We evaluate and the compare m-TSNE's visualizations with other competing approaches by conducting an user study. We show that m-TSNE's visualizations on the ATOM-HP dataset can extract patient activity level patterns and outliers (during chemotherapy cycle) which helps oncologists to study their treatment's effects on patient's fatigue. On the EEG dataset, we show that control and alcoholic subjects can be easily identified (separated) in our low-dimensional visualization which is not possible in the original high-dimensional space.

The organization of the paper is as follows: We discuss the related work in the field of data visualization in Section 2. Section 3 describes our m-TSNE visualization framework. Section 4 reports the empirical evaluation using the two aforementioned datasets and Section 5 concludes with discussion and future work.

2 Related Work

Due to the pervasiveness of time series data and its wide range of applications in healthcare monitoring, stock market analysis, traffic analysis, etc., understanding time-series data through visualization has attracted lots of attention, among which there have been efforts focusing on visualizing MTS data [8, 9, 10, 16]. These techniques visualize multivariate data by simply displaying individual variable data independently on one shared-space display [17] without treating multiple variables as a whole. Nevertheless, shared-space techniques may not be efficient to visualize high-dimensional MTS dataset since the resulting shared-space visualization may show a large amount of overlapping and clutter between different variables' time series, which is difficult to interpret. For instance, as shown in Figure 3, in [10], the authors proposed a technique named Multivariate Time Series Amalgam (MTSA), to jointly visualize multiple variables of MTS on a single display. Nevertheless, MTSA may not be suitable when the MTS dataset has a large number of measurements or has high dimensionality. The authors in [16] proposed a visualization method for multivariate data based on star chart [8] which represents each variable on an axis, with the axes arranged around a circle. An example of this technique is depicted in Figure 1. Another popular technique for visualizing MTS is parallel

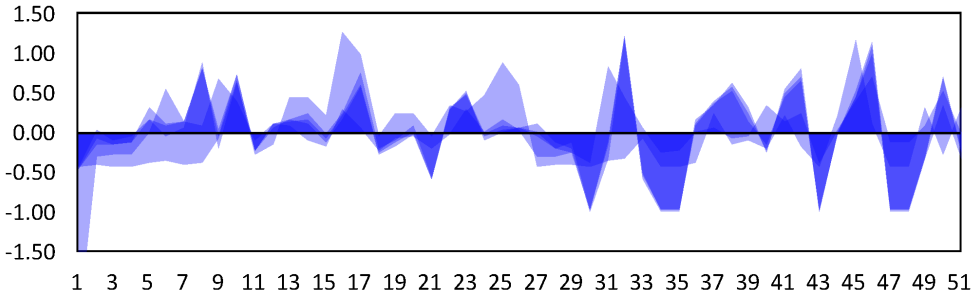

Figure 3: MTSA visualization [10] of the same raw data as Figure 1. x-axis represents the time instance (the date number), and y-axis represents the variables values.

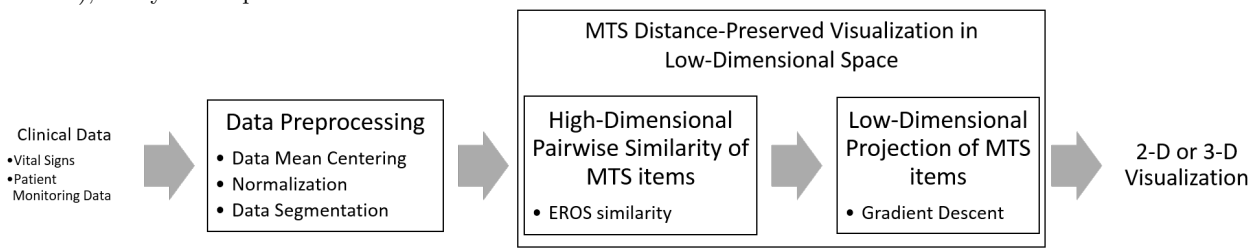

Figure 4: Pipeline of m-TSNE framework for MTS Visualization

coordinates [9] (shown in Figure 2). These methods have limitations in interpretability when the MTS has a large number of variables or a large number of time instances per variable. Unlike these studies, our proposed m-TSNE technique deals with multiple variables of MTS as a whole by applying dimensionality reduction techniques to project MTS data points to low-dimensional latent space and it also preserves the distances between the MTS in the original space.

3 MTS Visualization Framework

Figure 4 shows the pipeline of our proposed m-TSNE approach for MTS Visualization. First, the MTS data is processed by mean-centering and normalization; then it is segmented into multiple MTS items [12] (discussed in Section 3.1). EROS [6] is calculated to find the high-dimensional pairwise similarity between the MTS items (see Section 3.2). For visualization, the MTS items is projected to low-dimensional space with the guarantee that the high-dimensional pairwise similarity relation is preserved (explained in Section 3.3). That is, if two MTS items are neighbors (or far apart) in high-dimensional space, their low-dimensional projection points should also be close (or far apart) to each other. We use gradient descent method similar to [14] for dimensionality reduction and projection.

3.1 Data Preprocessing

Data preprocessing consists of data mean-centering, normalization, and segmentation. Mean-centering by subtracting the mean value of a variable, and normalization by dividing the variable with its standard deviation helps the MTS items to be in the same scale [12] which is needed for the further MTS similarity calculation step. The MTS data is also segmented into MTS items corresponding to the time instance (eg. an hour, a day). The trend of the data over the time should be visualized as the patterns of all MTS items.

3.2 MTS Pairwise High-dimensional Similarity

One of the widely used techniques for high-dimensional data similarity calculation is t-Distributed Stochastic Neighbor Embedding (t-SNE) [13]. In order to calculate the high-dimensional data points similarity, t-SNE

[13] computes the Euclidean distance between two data points, then converts the distance into conditional probabilities that represent similarities. The details of the techniques can be found in [18].

As our study concerns MTS, the similarity metric should be suitable for MTS data type and should also take into account the correlation between MTS variables. The Euclidean distance similarity metric is shown to be not suitable for time series data as acceleration and deceleration along the time axis is suboptimal for distance matching [19]. As a result, Dynamic Time Warping (DTW) is generally used to overcome the limitations of Euclidean distance metric [20]. However, both of these distance metrics only work for UTS data matching, as they do not consider the correlation between MTS variables. Therefore, the t-SNE technique may not perform well on MTS datasets if Euclidean distance or DTW is used as a similarity metric.

We propose to use EROS similarity metric [6] to overcome the shortcomings for above similarity metrics. EROS similarity metric is a technique based on the Principal Component Analysis method in which an MTS item is treated as a whole, i.e. it is not broken in multiple UTS, to preserve the correlation between variables. Let us denote MTS data by X, where $X = \{x_1, x_2, ..., x_k\}$ and x_i corresponds to a multivariate data point (high-dimensional) at time instance i. Each x_i is termed as an MTS item and it is represented as an m x n matrix, where m is the number of observations (eg. patients), and n is the number of variables. Each variable can be a vital sign variable such as heart rate, respiration rate, etc. or an activity monitoring variable such as step count, active hour, etc. Given 2 MTS items x_i and x_j, EROS first computes the eigenvectors and eigenvalues of each item. Thereafter, it measures the cosine similarities of the corresponding eigenvectors of x_i and x_j. Finally, EROS similarity is the weighted sum of all the cosine similarities of the eigenvectors. The weight is calculated as the aggregated value based on all the eigenvalues of the MTS items in the dataset. The EROS similarity metric is described in Equation 1:

$$EROS(x_i, x_j, w) = \sum_{l=1}^{n} w_l | < v_{il}, v_{jl} > | \tag{1}$$

Where, $EROS(x_i, x_j, w)$ is the similarity of MTS item x_i and MTS item x_j, $v_i = [v_{i1}, \cdots, v_{in}]$ and $v_j = [v_{j1}, \cdots, v_{jn}]$ are the two sets of eigenvectors of x_i and x_j respectively and w_l is the aggregated weight computed based on the eigenvalue corresponding to the l^{th} eigenvector in the weight vector w. The computed pairwise similarities of MTS items are used for projecting each MTS item into low-dimensional (2-D or 3-D) space. Section 3.3 describes the details of how to project MTS data points to lower dimensional space using a gradient descent method.

3.3 MTS Low-dimensional Projection

Principal Components Analysis (PCA) [12, 21] is a popular approach to map high-dimensional data into low-dimensional space. PCA is a linear mapping which focuses on preserving the low-dimensional projection of dissimilar data points far apart. However, PCA is shown not to be suitable for non-linear manifolds where preserving the low-dimensional projection of similar data points close to each other is important [13, 14]. To overcome this drawbacks of PCA, we propose to perform an optimization method: gradient descent (similar to the one used in t-SNE [14]) for m-TSNE where we minimize the mismatch between high-dimensional and low-dimensional spaces. The gradient descent method works by minimizing a cost function over all data points. The details of the cost function and performing gradient descent is explained in our technical report [18]. In the following section, we will discuss the empirical results of our m-TSNE approach, and compare it to PCA, Euclidean-based t-SNE, and DTW-based t-SNE approaches for MTS visualization.

4 Experiments

4.1 Datasets

We evaluate m-TSNE visualization using two healthcare datasets: ATOM-HP[1] dataset which is a coarse-grained time series dataset collected from our on-going study, and EEG dataset - a fine-grained time series dataset from UCI machine learning repository [22].

[1] ATOM-HP: Analytical Technologies to Objectively Measure Human Performance

ATOM-HP dataset : This dataset is collected to study how to quantify the activity levels of cancer patients undergoing chemotherapy treatment to complement the current clinical assessment. The patients suffer from treatment induced fatigue that affects their daily activities, which is usually measured by physicians using Eastern Cooperative Oncology Group (ECOG) scores [23]. However, this score tends to suffer from subjective bias. Therefore, a more robust objective measurement is needed to evaluate patients' activity / performance status, which is the goal of our study. In ATOM-HP, each patient carries a wearable sensor during their chemotherapy treatment cycle. The chemotherapy cycle consists of two chemotherapy visits. The date range between the two chemotherapy visits varies from two to three weeks. There are eight patients in our dataset (more patients are being enrolled in this on-going study). The daily data of each patient is an MTS which has five variables: number of steps, total calories, average heart rate, peak heart rate, and lowest heart rate. Collected data is sampled every hour and for each patient, data is collected for at least 50 consecutive days.

EEG dataset [15] : The dataset is collected to examine if EEG correlates genetic predisposition to alcoholism. There are three versions of the dataset. In our work, we use the large dataset version that has 10 control subjects, and 10 alcoholic subjects. Each subject performs 30 trials which can be classified as three trial types: exposure to a single stimulus, exposure to two matching stimulus, and exposure to two non-matching stimulus [24]. In total there are 600 (=20x30) trials. The data of one trial is an MTS of 64 variables corresponding to 64 electrodes placed on the subject's scalps. The data sample rate is 256Hz.

4.2 Experimental Setup

To evaluate our framework, we compare m-TSNE to the visualizations of PCA, Euclidean-based t-SNE, and DTW-based t-SNE methods. In the ATOM-HP dataset, given an MTS data of one subject, we are interested in visualization of MTS to show the trends and outliers in the subject's daily activity performance during the chemotherapy treatment. The subject's MTS data is represented as an m x 5 matrix with m is the total hours of monitoring, and 5 is the number of variables (Section 4.1). For instance, the size of the MTS matrix data of the subject in Figure 5 is 1224 x 5. As the study considers monitoring daily performance, the MTS is segmented into multiple MTS items of 24 x 5 matrices which represent the 24 hours of the data after mean-centering, and normalization. In Figure 5, we have a total of 51 (= 1224/24) data points corresponding to 51 MTS items data matrices collected during 51 days when the subject was involved in the study. For m-TSNE, we calculate the pairwise similarity using EROS similarity metric (Equation 1). The pairwise similarities are put through the gradient descent to compute the MTS items' projection in low-dimensional (2-D or 3-D) space (Figure 5, Figure 7). For the EEG dataset, each subject's trial is represented as an MTS matrix of 256 x 64 where 256 is the number of observations and 64 is the number of EEG variables. The pairwise EROS similarities of 600 MTS items corresponding to 600 trials is calculated. The similarities are then used to compute the projection of 600 MTS as Figure 6.

PCA: The MTS is preprocessed, and segmented into MTS items as above (Section 4.2). After preprocessing, the daily aggregated multivariate data points are computed as sum of number of steps, sum of calories, average heart rate, peak heart rate, and lowest heart rate over 24 hours. All aggregated data of MTS items are put together as a matrix of m' x 5 where 5 is the number of variables, and m' is the number of data points with each row representing an aggregated multivariate data point. PCA is used for low-dimensional projection and visualization as shown in Figure 8.

Euclidean-based t-SNE: After data preprocessing, a pairwise Euclidean distance for each data points pair is calculated. Based on the Euclidean distance, a pairwise probability similarity matrix of each pair of data points is computed as in [18].

DTW-based t-SNE: Since Euclidean distance does not work well for time series data, we can measure MTS similarity using DTW distance metric. This approach is similar to Euclidean-based t-SNE, but uses the pairwise Dynamic Time Warping distance instead of Euclidean distance.

m-TSNE: m-TSNE calculates EROS pairwise similarity matrix, and computes MTS projection using the gradient descent method as described in Section 3.

We implemented above four methods in Python 2.7.11 and we will release our code on GitHub.

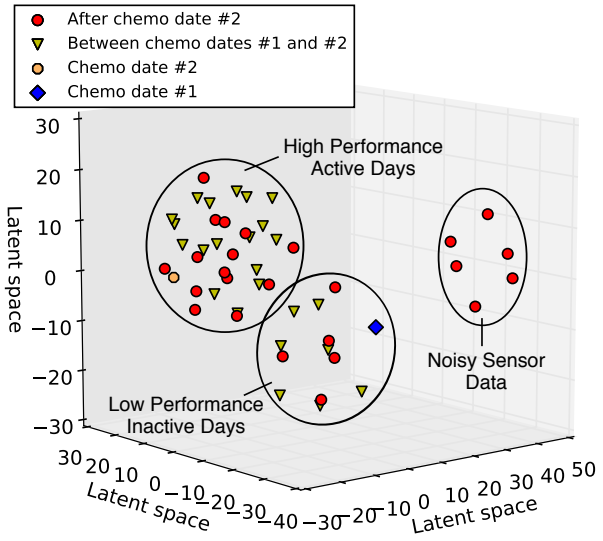

Figure 5: m-TSNE visualization of ATOM-HP dataset.

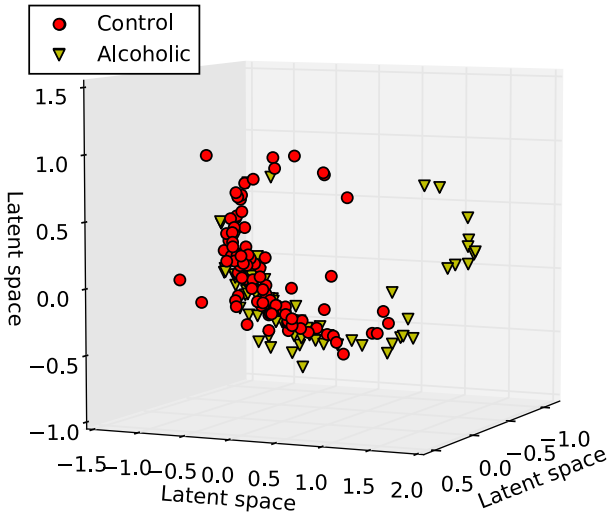

Figure 6: m-TSNE visualization of EEG dataset.

4.3 Visualization Results

Figure 5 shows that the visualization of the MTS items of a subject undergoing treatment using m-TSNE. In this figure, there are 51 points corresponding to 51 days of chemotherapy treatment cycle. Each data point is labeled by the date order whether it is before or after a chemotherapy date in the study cycle. As can be seen, the figure shows that the 51 points form 3 clusters. To understand the three clusters, a 2-D display with annotation for each data point is provided in Figure 7. The annotation are in a format [**date number**]_[**step counts**] (step counts variable is chosen as it is easier to interpret human performance using this variable as compared to other variables). The figure was also shown to a health professional. Based on the health professionals' inputs, and the variables values, in Figure 5, these clusters can be interpreted as: (1) cluster of high performance / active days (the left-most cluster with a high number of step counts for each data point (more than 800)), (2) cluster of low performance / inactive days (the lower cluster with a low number of step counts for each data point), (3) cluster of noisy data (outliers) from the sensors (the right-most cluster) (noisy data appears to have abnormal values due to sensors). This is one of the insights obtained from our visualization approach that can help the healthcare professionals to understand the effect of the treatment using data collected from wearable sensors. Moreover, the figure also shows that the subject has low activities for a few days immediately following the two chemotherapy dates. This insight indicates that the subject may suffer from fatigue due to the chemotherapy session. We believe that these insights are quite helpful for oncologists to study their patient's activity performance, and also these might help them in designing better objective measures to quantify human performance.

For comparison, we provide the visualizations of the same subject using PCA, Euclidean-based t-SNE and DTW-based t-SNE methods in Figure 8, Figure 9, and Figure 10, respectively. Note, the star chart, parallel coordinates chart, and MTSA of the same subject's data are also shown in Figure 1, Figure 2, and Figure 3 respectively. It is clear that these visualizations do not provide clear insights about the patient's activities / performance status. Euclidean-based, and DTW-based t-SNE figures (Figure 9, Figure 10) show that the noisy sensor data points are close to each other. However, they do not give clear distinct clusters like m-TSNE. Based on these figures, it may be difficult to interpret and understand the data insights as these approaches show overlapping clusters which appear as a cloud of data points. PCA (Figure 8) might provide some insights of the subject performance along the axis of highest principle component, however as shown in the figure it does not form distinct clusters for outliers or for different levels of activity.

Figure 6 provides m-TSNE visualization for the EEG dataset. Each data point is one subject's MTS data performing a trial, and is labeled based on the subject category: *control* and *alcoholic*. The figure clearly depicts a manifold in which the points representing the control group lie inside, and are covered by alcoholic group. It also shows that all the outliers in our visualization belong to alcoholic group, providing

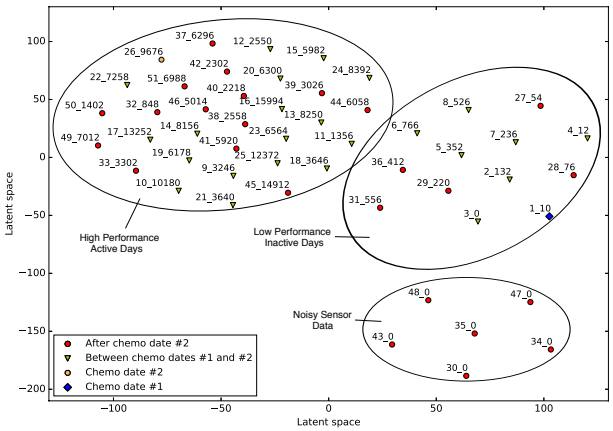

Figure 7: m-TSNE 2-D visualization of ATOM-HP dataset with annotations.

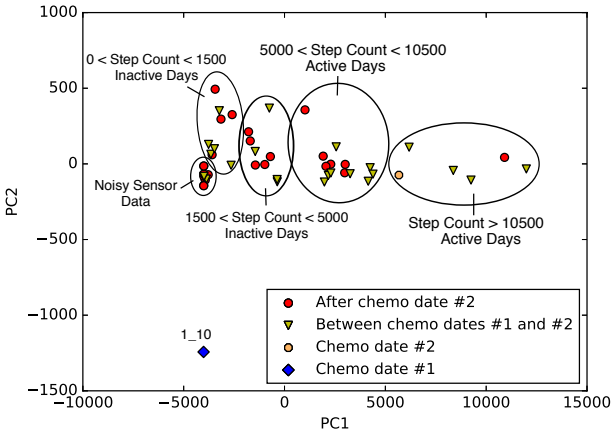

Figure 8: PCA 2-D visualization of ATOM-HP dataset.

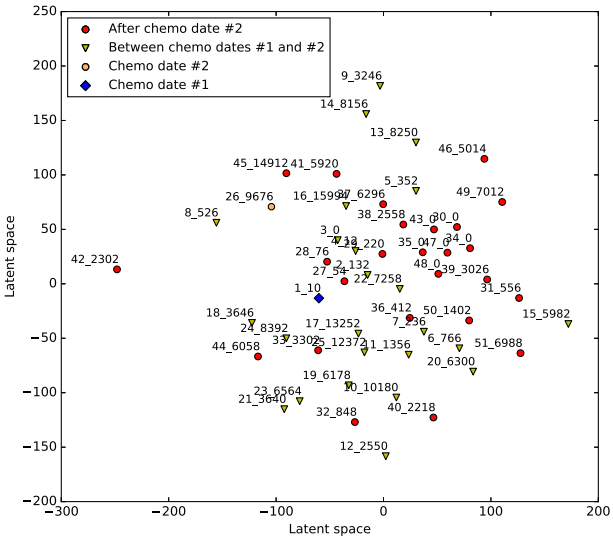

Figure 9: Euclidean-based t-SNE 2-D visualization of ATOM-HP dataset with annotations.

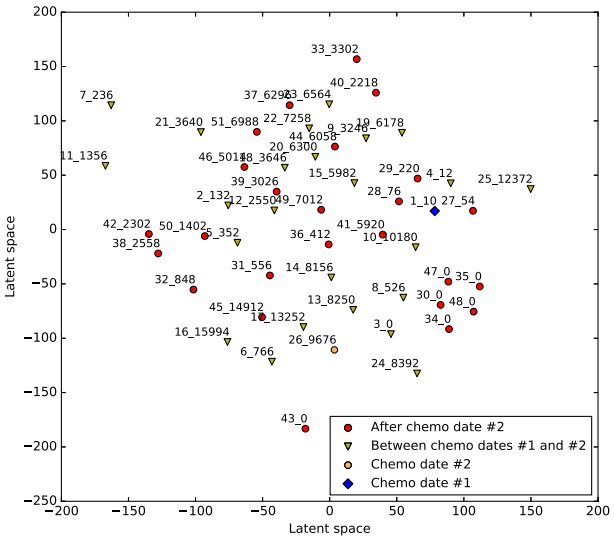

Figure 10: DTW-based t-SNE 2-D visualization of ATOM-HP dataset with annotations.

interpretable insights which are not extracted by PCA or the other MTS visualization techniques.

User Study: We evaluated the interpretability of three techniques: PCA, DTW-based t-SNE and m-TSNE by conducting a controlled user study with 6 non-healthcare professionals. We removed the names of the techniques, labels and color codes from the visualization results to avoid bias in users' interpretability. Each user was shown the visualizations of these three techniques on all subjects in ATOM-HP dataset and instructed to assign a score if they could find (interpretable) clusters, trends or outliers. For each technique, an user assigned a score (1 - the lowest score, 2 or 3 - the best score) in such a way that better techniques received higher scores. We aggregated the scores over all users and report out findings here: m-TSNE obtained the highest score of 2.48, PCA obtained a score of 1.92, and DTW-based t-SNE obtained the lowest score of 1.6. An oncologist was also included in our user study to verify our visualizations and insights. He agreed that the clusters found by our approach is very useful to study the patients' fatigue during their treatment cycle. This user study shows that m-TSNE could provide interpretability and insights when compared to the other competing methods.

5 Conclusion and Future Work

In this paper, we proposed m-TSNE: a framework to visualize high-dimensional MTS data. m-TSNE uses EROS to compute similarity between MTS data points and projects them to low-dimensional space for visualization. Empirical evaluation on two healthcare datasets showed that our approach provides interpretable insights via visualization while the other visualization methods which use PCA, Euclidean-based t-SNE, and DTW-based t-SNE, are more difficult to interpret. These insights could help healthcare professionals to evaluate their patients' performance. For future work, we plan to extend our work by building a tool for showing the visualization of the MTS dynamically.

Acknowledgments: This research has been funded in part by the National Cancer Institute (award number P30CA014089), National Institutes of Health, Department of Health and Human Services, under Contract No. HHSN261201500003B, and the USC Integrated Media Systems Center. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of the sponsors.

References

[1] W. Raghupathi and V. Raghupathi, "Big data analytics in healthcare: promise and potential," *Health Information Science and Systems*, vol. 2, no. 1, p. 1, 2014.

[2] J. J. Van Wijk and E. R. Van Selow, "Cluster and calendar based visualization of time series data," in *Information Visualization, 1999.(Info Vis' 99) Proceedings. 1999 IEEE Symposium on*. IEEE, 1999, pp. 4–9.

[3] A. Aris, B. Shneiderman, C. Plaisant, G. Shmueli, and W. Jank, "Representing unevenly-spaced time series data for visualization and interactive exploration," in *IFIP Conference on Human-Computer Interaction*. Springer, 2005.

[4] Y. Li, J. Lin, and T. Oates, "Visualizing variable-length time series motifs." in *SDM*. SIAM, 2012, pp. 895–906.

[5] C. Shi, W. Cui, S. Liu, P. Xu, W. Chen, and H. Qu, "Rankexplorer: Visualization of ranking changes in large time series data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2669–2678, 2012.

[6] K. Yang and C. Shahabi, "A pca-based similarity measure for multivariate time series," in *Proceedings of the 2nd ACM international workshop on Multimedia databases*. ACM, 2004, pp. 65–74.

[7] J. Heer, M. Bostock, and V. Ogievetsky, "A tour through the visualization zoo. retrieved may 17, 2013," 2010.

[8] J. M. Chambers, *Graphical methods for data analysis*, 1983.

[9] A. Inselberg, "The plane with parallel coordinates," *The visual computer*, vol. 1, no. 2, pp. 69–91, 1985.

[10] C. U. Lehmann and J. Fackler, "Visualizing multivariate time series data to detect specific medical conditions," 2008.

[11] V. West, D. Borland, and W. E. Hammond, "Visualization of ehr and health related data for information discovery," in *Workshop on Visual Analytics in Healthcare*, 2013.

[12] I. Jolliffe, *Principal component analysis*. Wiley Online Library, 2002.

[13] G. E. Hinton and S. T. Roweis, "Stochastic neighbor embedding," in *Advances in neural information processing systems*, 2002, pp. 833–840.

[14] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, no. Nov, pp. 2579–2605, 2008.

[15] X. L. Zhang, H. Begleiter, B. Porjesz, W. Wang, and A. Litke, "Event related potentials during object recognition tasks," *Brain Research Bulletin*, vol. 38, no. 6, pp. 531–538, 1995.

[16] D. Borland, V. West, and W. Hammond, "Multivariate visualization of system-wide national health service data using radial coordinates," in *Proc. Workshop on Visual Analytics in Healthcare*, 2014.

[17] W. Javed, B. McDonnel, and N. Elmqvist, "Graphical perception of multiple time series," *IEEE transactions on visualization and computer graphics*, vol. 16, no. 6, pp. 927–934, 2010.

[18] M. Nguyen, S. Purushotham, H. To, and C. Shahabi, "m-tsne: A framework for visualizing high-dimensional multivariate time series (technical report)," *University of Southern California, 2016. http://infolab.usc.edu/DocsDemos/mtsne.pdf*.

[19] O. Maimon and L. Rokach, *Data mining and knowledge discovery handbook*. Springer, 2005, vol. 2.

[20] M. Müller, "Dynamic time warping," *Information retrieval for music and motion*, pp. 69–84, 2007.

[21] W. Krzanowski, "Between-groups comparison of principal components," *Journal of the American Statistical Association*, vol. 74, no. 367, pp. 703–707, 1979.

[22] M. Lichman, "UCI machine learning repository," 2013. [Online]. Available: http://archive.ics.uci.edu/ml

[23] M. M. Oken, R. H. Creech, D. C. Tormey, J. Horton, T. E. Davis, E. T. Mcfadden, and P. P. Carbone, "Toxicity and response criteria of the eastern cooperative oncology group." *American journal of clinical oncology*, vol. 5, no. 6, 1982.

[24] J. G. Snodgrass and M. Vanderwart, "A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity." *Journal of experimental psychology: Human learning and memory*, vol. 6, no. 2, 1980.

Interactive Visualization and Exploration of Patient Progression in a Hospital Setting

Wathsala Widanagamaachchi* SCI Institute, University of Utah
Yarden Livnat† SCI Institute, University of Utah
Peer-Timo Bremer‡ Lawrence Livermore National Laboratory
Scott Duvall§ Internal Medicine, University of Utah
Valerio Pascucci¶ SCI Institute, University of Utah

ABSTRACT

As medical organizations increasingly adopt the use of electronic health records (EHRs), large volumes of clinical data are being captured on a daily basis. These data provide comprehensive information about patients and have the potential to improve a wide range of application domains in healthcare. Physicians and clinical researchers are interested in finding effective ways to understand this abundance of data. Use of visual analytics to analyze and explore healthcare data is one such research direction. In this work, we present a visualization and analysis environment to understand patient progression over time. Through the use of optimized data structures and progressive visualization techniques, we allow users to interactively explore how patients and their progression change over time. Compared to existing techniques, our work provides additional flexibility in analyzing patient data and has the potential to be used in a real-time hospital setting. Finally, we demonstrate the utility of our approach using a publicly available intensive care unit (ICU) database.

1 INTRODUCTION

The US healthcare system is producing hundreds of thousands of patient records detailing a wide range of information from admission times and dates, to symptoms and outcomes. Until recently, this data has been difficult to access, especially in bulk, often lacked a useful organization, and thus has been generally underutilized for clinical research. With the increasing use of EHRs, this paradigm changes, allowing researchers easy access to a large collection of information. If used effectively, this data may lead to better predictions of patient outcomes, personalized medication, and more targeted interventions. However, to realize this potential requires the ability to understand the clinical data in detail. Given the massive amounts of available data, for example, ICUs may collect real-time data streams of all patients [1], which implies automatic or semiautomatic techniques to identify and explore interesting patterns and underlying trends. In this context, visualizing and exploring patient progression over time can provide valuable insights and facilitate the decision-making of physicians and clinical researchers.

Several factors need to be taken into consideration when analyzing this type of data: First, given the large number of patients, an individual, per-patient analysis is time-consuming and does not lend itself to finding commonalities and trends. Instead, patients should be grouped according to various criteria, such as symptoms, outcomes, etc. Second, to compare groups of patients who arrive at different times, their records must be aligned, for example, by their time of admission, time of major procedures or other common factors. Third, patient progression over time needs to be presented in a concise manner to allow simultaneous exploration of large numbers of patients. Finally, to utilize such a system in a hospital setting, the analysis must be interactive, allowing users to quickly explore different hypotheses.

The ideal system described above presents a number of practical challenges, especially for the large databases of interest. First, there exist a number of potentially interesting metrics by which to group patients and thus any analysis must be flexible and efficient enough to change the metric on-the-fly. Furthermore, whereas some metrics are easy to apply and absolute (e.g., splits by gender), others depend on specifying a *similarity threshold* that determines when two patients are considered to be in the same group. However, in practice this threshold is typically not known a priori, and in fact understanding how patient distributions and progression change with different thresholds may provide important insights. Most existing approaches focus on a single metric and a preselected threshold [37, 38]; we present a system that allows users to freely explore metrics and thresholds in an interactive setting.

Another challenge is the size and complexity of the data. Given a large number of patients and high temporal resolution, it is often difficult to grasp the progression of certain groups, let alone identify salient ones. Therefore, presenting data in a concise manner and providing support for various parameter selections and simplifications is crucial to provide the necessary insights.

From an analysis perspective, providing an effective exploration of patient progression requires three abilities: first, grouping patients within a time step at different similarity thresholds; second, correlating patient groups over time; and third, interactively visualizing and exploring patient progression to understand how different similarities affect their behavior. In this paper, we focus on extracting patient groups across multiple patient similarities and exploring their progression with the aid of tracking graphs, where a concise representation of feature evolution is captured as a collection of feature tracks, see Figure 1. We provide clinical researchers with a visualization and analysis environment that is developed based on our earlier research in the scientific domain [35, 36]. This prior system couples feature grouping and correlation components with visualization techniques to explore the temporal evolution of features in combustion data sets.

Our visual analytics process data in several steps. First, we use the patient similarity metric introduced in [18] to group patients across multiple similarity thresholds. Then, patient groups are correlated over time by tracking the individual patients within. In order to allow interactive extraction of data, our system uses optimized data structures to store these patient group and correlation details. Within the system, tracking graphs are used to present a global concise view of patient progression, and progressive visualization techniques are employed to enable interactive exploration of data. Finally, in collaboration with clinical researchers, we apply our visualization and analysis environment to a publicly available ICU

*e-mail:wathsy@sci.utah.edu
†e-mail:yarden@sci.utah.edu
‡e-mail:bremer5@llnl.gov
§e-mail:scott.duvall@utah.edu
¶e-mail:pascucci@sci.utah.edu

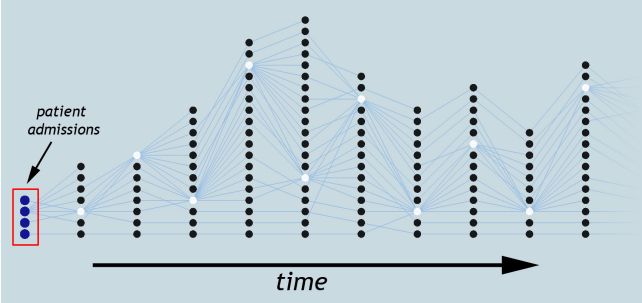

Figure 1: An example tracking graph showing patient progression over time. Each node represents a patient group and its "track" shows how that group progresses over time.

database, the clinical database of Multiparameter Intelligent Monitoring in Intensive Care (MIMIC II) databases [30], and explore the temporal progression of patients for varying similarity thresholds.

2 Related Work

A subset of the relevant related work is presented here to provide context and background for our research work. Analyzing time-varying data sets usually involves feature extraction and tracking steps. For healthcare data, tracking the progression of patient groups, i.e., the features-of-interest, is relevant to clinical researchers. Among the many feature definitions and their computation techniques can be found in the literature, techniques that extract feature information for all or a large range of values in a single pass are particularly useful. These techniques often result in hierarchical representations. For instance, hierarchical clustering [34, 15, 39, 3, 6] and various other topological techniques [4, 32, 5] have been used to effectively capture flexible feature hierarchies.

Hierarchical clustering is considered to be one of the most popular methods for creating a feature hierarchy. It partitions data into homogeneous groups based on a measure of similarity through the use of clustering. Depending on the similarity measure used, the results can lead to very different hierarchies. Moreover, many sequential and parallel algorithms for hierarchical clustering are available in the literature. Several important results on sequential algorithms are presented in [21, 10] and details on previous parallel algorithms for hierarchical clustering are summarized in [22]. This type of clustering imposes a hierarchical structure on the underlying data irrespective of whether such a structure is appropriate. However, due to its simplicity, many applications have used this method to explore the clustering hierarchy of features. In this work, we also make use of hierarchical clustering to group patients within a time step at different similarity thresholds.

In topological analysis, techniques exist that are able to efficiently extract and encode entire feature families in a single analysis pass. Reeb graphs [24], contour trees [7], merge trees [5], and Morse-Smale complexes [4] are several such techniques. Among them, Reeb graph, contour tree, and merge tree are contour-based and the Morse-Smale complex is gradient-based. As a result, the Morse-Smale complex captures very different structural information.

Visualizing the temporal evolution of features has long been a problem of interest within the visualization community. Depending on the subject area, many different techniques have been developed to address this problem. Traditionally, abstraction, illustration, morphing or animation-based techniques [20, 14, 16] have been used to visualize temporal evolution of features. Other techniques such as change detection [31] and high-dimensional projection [17] have also been used in the past. Tracking graphs that show the feature evolution as a collection of feature tracks that split or merge over time are considered to be an effective representation for visualizing feature evolution [28, 35]. These graphs provide concise global views of feature evolution and are more amenable to filtering and simplifications. As clinical researchers are particularly interested in concise representations, we make use of tracking graphs to visualize patient progression over time.

[29] includes a comprehensive survey of information visualization systems used to visualize, explore, and query EHRs can be found. These approaches related to EHRs can be broadly categorized into two categories: those that focus on a single patient record [27, 13, 26] and those concerned with a collection of patient records [33, 37, 38, 19]. Approaches in the first category focus on providing comprehensive information about a single patient (e.g., patient history, significant events, medication, and treatment), and the second category aims at presenting an overview from multiple patients. The latter provides less detail on each individual patient and focuses more on recognizing patterns and outliers within patient groups. Among these approaches that fall in the second category, LifeFlow [38] and OutFlow [37] are particularly interesting as they visualize event sequences in EHRs. LifeFlow uses color for a compact view and OutFlow uses a graph-based representation. In contrast, we do not visualize the progression of patient groups as an event. At a particular time step, the current event of a patient is one of the parameters considered within the patient similarity metric used. Also, within our system any similarity metric can be used to define patient similarities, providing more flexibility.

3 System Components

An interactive visualization and analysis environment is essential to gain an in-depth understanding of patient progression. In this work, we refine a prior system that relies on dynamically constructed tracking graphs to enable feature extraction, tracking, and simplification [35, 36]. Although this system is designed to study general time-varying features, so far it has been applied to analyze features only in scientific simulations. We extend its functionality to effectively visualize patient progression in healthcare data. This section provides a comprehensive description of our system partitioned into several subsections dealing with: patient grouping, patient correlation, visualization, exploration, and implementation.

3.1 Grouping Patients Within a Time Step

The first step towards understanding clinical data is defining its features and a time step size based on which subsequent analysis is to be conducted. For our intended research, the feature-of-interest is a patient group (i.e., similar set of patients), and a day is considered to be the appropriate time step size. Next, for each time step in the data set, these patient groups need to be extracted and aligned. In this work, to ensure all patients' hospital stays start at the same time, we align data based on a patient's admission time.

Once features are extracted and aligned, they should be grouped based on a appropriate grouping method. By maintaining a notion of scale, this feature grouping naturally approximates a meaningful hierarchy. The naive approach of creating this hierarchy is by exhaustively precomputing all possible features at all possible scales. Many popular grouping algorithms also produce nested sets of features for varying scale, which in turn creates these types of hierarchies (e.g., hierarchical clustering techniques progressively merge elements [8] and threshold-based segmentation creates increasingly larger regions [4]). In our case, we use hierarchical clustering. For each time step in the data set, patient groups are clustered based on their similarity to generate a hierarchical representation in the form of a tree. During clustering, we use the metric of [18] to define patient similarities but any other similarity metric could be used as well.

Figure 2(a) shows an example where such a hierarchy is constructed by progressively merging individual patients, with the most similar ones clustered first. Each leaf in the hierarchy represents a patient and each branch a patient group. Along with the hierarchy, various patient group-based attributes such as patient count, mean age and mean heart rate are computed and stored on a per-branch basis. For a given data set, an offline preprocessing step is used to compute these patient hierarchies, and the results are stored in a look-up structure to allow interactive exploration of patient groups. Within this look-up structure, for each patient group, its parent details and patient group-based attributes are stored. To determine correspondences across patients later, each patient is marked with a unique ID. This hierarchy is computed for each time step in the data, and is stored in a separate file to allow interactive exploration of patient groups.

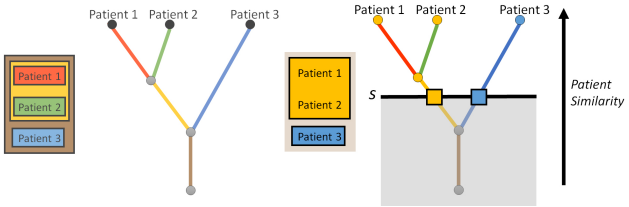

(a) (b)

Figure 2: (a) A patient hierarchy constructed by progressively merging similar patients, with the most similar ones clustered first. (b) To extract patient groups, the hierarchy is cut at a fixed threshold, resulting in a forest of subtrees, where each subtree represents a patient group.

Once the hierarchy is computed, patient groups and their attributes can be quickly and easily extracted for any similarity threshold within its range, see Figure 2(b). Given a similarity threshold s within the full range of r, the corresponding patient groups can be extracted by "cutting" the hierarchy at s. This creates a forest of subtrees, where each subtree represents a patient group existing at s.

3.2 Correlating Groups of Patients Over Time

Once patient groups are identified, the next step is to correlate them over time by tracking individual patients within. Two patient groups in consecutive time steps are considered to be correlated if they share at least one patient. All such correlations are extracted for each time step. To efficiently store and interactively extract these correlation details, we utilize the meta-graph structure of [35]. Similar to the aforementioned patient hierarchy, this meta-graph structure is able to encode patient group correlations and their attributes for a range of similarity thresholds.

The meta-graph is generated in two steps. First, per-patient correlations are computed using the patient IDs computed above. For example, two patients in consecutive time steps are considered to be correlated if they have the same ID. As individual patients are represented by leaf branches in the feature hierarchy, this step results in correlations across leaf branches in consecutive time steps. If a correlation exists, we assign an edge with the weight of 1 across the two corresponding leaf branches, $(b_i^t, b_j^{t+1}, 1)$. Second, these per-patient correlations are accumulated along the feature hierarchy to compute the per-patient group correlations. At the accumulation time, if a correspondence already exists, we accumulate only the edge weights.

Just as in the patient hierarchy, various correlation-based attributes such as the amount of patient overlap are computed and stored within the meta-graph structure. Again, once the meta-graph is computed, patient group correlations and their attributes can be quickly extracted for any similarity threshold within the full parameter range. For a selected similarity threshold s, first, patient groups existing at s for each time step in the data set are obtained using the precomputed patient hierarchies. Then, correlations that exist across those extracted patient groups are obtained from the meta-graph structure. Together, these extracted patient groups and their correlations form the tracking graph at f, see Figure 3. This meta-graph structure is also created in an offline preprocessing step and the resulting structure is stored in multiple files (i.e., one file per time step), each containing a set of edges representing its correlations to patient groups in the next time step.

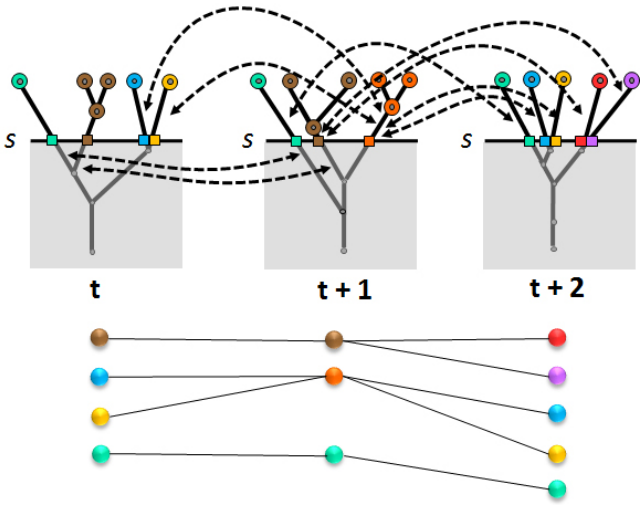

Figure 3: Tracking graph construction. For a similarity threshold s, first, patient groups existing at that value are obtained from corresponding patient hierarchies. Then, the meta-graph is used to extract correlation details. Here, the correlations extracted are indicated with black arrows. The resulting tracking graph is displayed at the bottom.

3.3 Visualizing and Exploring Patient Progression

Our system for exploring patient progression over time contains three views: patient grouping view, patient progression view, and patient view, see Figure 4. Within each view, various progressive visualization techniques are employed to achieve interactivity. For instance, data is always presented with respect to a focus time step that is processed first. Data for the neighboring time steps is then extracted and presented in order of increasing distance. All views designed for only a single time step (i.e., patient grouping view and patient view), use the focus to determine their time step. The parameters such as hierarchy parameters and other filter parameters are coordinated across all views to provide a fully linked analysis environment.

3.3.1 Patient Grouping View

To enable researchers to gain a quick visual understanding of how patients group together for varying similarity thresholds, the patient hierarchy of the focus time step is visualized within this view. As the similarity threshold is changed, active patient groups within the hierarchy are also highlighted. In Figure 4(b), the selected similarity threshold within the hierarchy is displayed in a brown vertical line, and the active patient groups are highlighted in prominent colors.

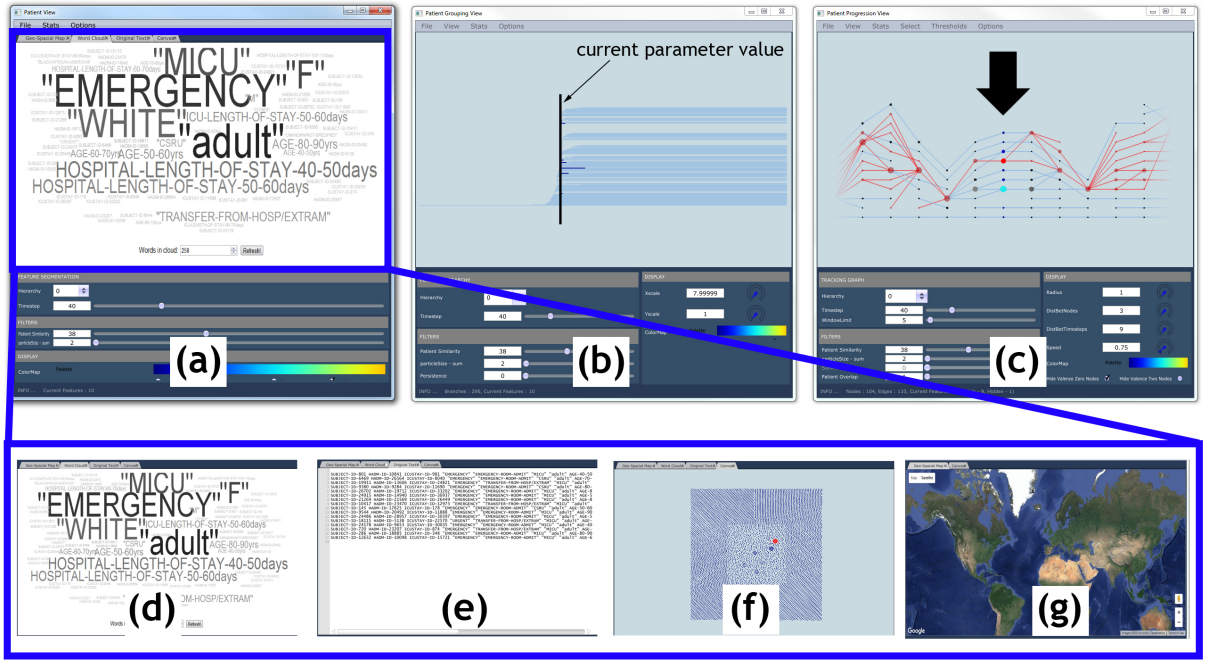

Figure 4: Our system contains three views: (a) patient view, (b) patient grouping view, and (c) patient progression view. The patient view consists of several subcomponents: (d) a word cloud, (e) textual, (f) geometric embedding, and (g) geospatial views. The patient grouping view shows the hierarchy for the focus time step and the patient progression view displays the tracking graph for the current focus and time window. Within the patient progression view, nodes are scaled based on the patient group size, and the focus time step is indicated with a black arrow. Here, a patient group is selected, which results in its progression being highlighted (indicated in red within the patient progression view). In the patient view, the selected patient group's details are displayed.

3.3.2 Patient Progression View

This view visualizes the temporal progression of patients using tracking graphs. Starting from the user-defined focus time step, nodes and edges are iteratively added both forward and backward in time up to the user-defined time window to create the tracking graph, see Figure 4(c). Each node in the graph represents a patient group. A set of nodes in the same x coordinate indicates groups in one time step and edges across them indicate their correlations. For visual clarity, nodes in the focus time step are always displayed in prominent colors. Progressive techniques as in [35], specifically, a fast initial graph layout and a slower greedy one, are used to visualize these tracking graphs.

3.3.3 Patient View

Several visualization techniques are combined here to present a specialized view of patients. Specifically, we integrate word cloud, textual, geometric embeddings, and geospatial visualizations, see Figure 4(a).

- Word Cloud Visualization
  This component is dedicated to providing a quick overview of textual information regarding patients. For a selected patient group, a word cloud is constructed from the patient group-based attributes stored within the patient hierarchy, see Figure 4(d). Here, to obtain more intuitive overviews, the numerical attributes are converted into ranges. This visualization displays high-frequency words using bigger fonts and brighter colors, and others in faded and smaller fonts.

- Textual Visualization
  As the name suggests, this visualizes textual details of patients in their native domain (i.e., as text), see Figure 4(e). For a selected patient group, textual visualization displays its attributes such as hospital admission ID, patient ID, care unit and age.

- Geometric Embedding Visualization
  Regardless of the data type, visualizing geometric embedding reveals interesting details and trends about data. This view visualizes the geometric embedding of patients in either 2D or 3D, see Figure 4(f). As geometric embedding details are not very obvious for the clinical data, for each time step, the GraphViz [11] 'neato' layout algorithm together with patient similarity details is used to compute the 2D embedding of patients.

- Geospatial Visualization
  When relevant information is available, we allow data exploration to be augmented with geospatial visualizations, see Figure 4(g). For instance, if a patient has his physical location details available for each moment in time (both during and/or prior to his hospital stay), this information will be visualized within this view. In addition to visualizing patient geospatial locations, their trajectories can also be displayed to easily identify data trends related to geographic locations.

3.3.4 Interactive Exploration

As tracking graphs can easily become complex and difficult to understand, various simplifications have to be performed on them to successfully understand their underlying trends. Specifically, we enable several simplification options. Through the linked-view interface, researchers are allowed to explore data sets by changing the focus time step and time window. They can select a particular day within a patient's hospital stay, expand and contract its neighboring days to view progression both forward and backward in time.

Within our system, the similarity threshold within the patient hierarchy, correlation amount within the meta-graph and other attribute values available (i.e., patient group-based and correlation-based), can all be explored. We also allow tracking graphs to be filtered by the length of stay of a patient group, which enables small patient stays to be eliminated from the analysis.

Valence two and zero nodes of a tracking graph can be hidden to prevent visual clutter, nodes can be scaled based on their size, and progressions of certain patient groups can be highlighted. To help researchers maintain context across systems' views, we also make use of correlated color maps and allow nodes to be colored using various patient group-based attributes. All these options combined enable researchers to interactively simplify tracking graphs, isolate interesting patient progressions and explore their parameter space.

3.4 Implementation

Our system is implemented using the ViSUS framework [25, 23], which provides the basic building blocks for designing a streaming, asynchronous dataflow. Figure 5 shows the dataflow utilized within our system.

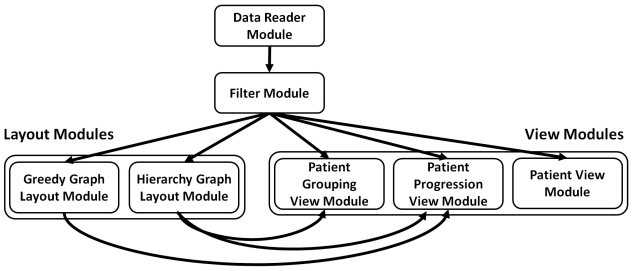

Figure 5: Our system's dataflow contains several modules. The data is read into the system using the `Data Reader` module. Then, this node and edge information is filtered according to the current parameters within the `Filter` module. This module sends the resultant data to `Layout` and `View` modules simultaneously. The two `Layout` modules within the dataflow compute the relevant graph layouts and the `View` modules render the information received.

The *Data Reader* module is dedicated to reading data into the system. It checks whether all data required for the current tracking graph has been loaded. If needed, it loads the required data and passes it to the *Filter* module. This module filters the received patient group and correlation details for the current parameter and attributes values. This filtered information is then simultaneously sent to layout and view modules.

Each of the two layout modules computes a graph layout and sends those layout details to the relevant view modules for rendering. The *Hierarchy Graph Layout* module computes the initial layout for the tracking graph and sends this information to the *Patient Grouping View* and *Patient Progression View* modules. This hierarchy graph layout is computed only once for each time step as the data is read for the first time. The second layout module, *Greedy Graph Layout*, computes a greedy layout for the tracking graph each time its parameters change and passes them to the *Patient Progression View* module. This greedy layout is computed to make sure the edge crossings are minimized within the tracking graph.

Our dataflow contains three view modules. The first view module, *Patient Grouping View*, visualizes the patient hierarchy of the focus time step. Once the module receives the necessary node and hierarchy details from *Filter* module and the layout details from *Hierarchy Graph Layout* module, it renders the patient hierarchy. The *Patient Progression View* module initially renders the tracking graph using the hierarchy graph layout. Then, as the greedy layout becomes available, it is integrated with the current graph. The third view module, *Patient View*, provides more specific views of patients (geometric embedding, geospatial, word cloud, and textual visualizations). Once this module receives the required data, depending on which visualization mode is selected, the corresponding computations and renderings are triggered. Each time parameters and/or selections are changed, the current processing within the dataflow is interrupted and restarted. However, rendering within the views maintains the current state for visual continuity.

Most parts of our system (except word cloud and geospatial visualizations) are implemented in C++ and use OpenGL rendering. The word cloud and geospatial visualizations make use of JavaScript libraries and functions such as d3-cloud [9], a Wordle-inspired word cloud layout, and Google maps [12]. Within the system, the integration between C++ and JavaScript is achieved using the Awesomium library [2], which enables C++ code to be seamlessly integrated with HTML UI and to maintain interactions across the two.

General	Patient demographics, hospital admissions, discharge dates, room tracking, death dates (in or out of the hospital), ICD-9 codes, unique code for healthcare provider, and type (RN, MD, RT, etc).
Physiological	Hourly vital sign metrics, SAPS, SOFA, ventilator settings, etc.
Medications	IV meds, provider order entry data, etc.
Lab Tests	Chemistry, hematology, ABGs, imaging, etc.
Fluid Balance	Intake (solutions, blood, etc), output (urine, estimated blood loss, etc).
Notes & Reports	Discharge summary, nursing progress notes, etc; cardiac catheterization, ECG, radiology, and echo reports.

Table 1: An overview of the data categories within MIMIC II clinical database

4 Results

We enable clinical researchers to study the progression of patients via interactive exploration of dynamically constructed tracking graphs. The effectiveness of our framework is demonstrated with the use of a publicly available ICU database.

The clinical database of Multiparameter Intelligent Monitoring in Intensive Care (MIMIC II) databases [30] contains comprehensive EHR data collected from hospital medical information systems (both patient bedside workstations and hospital archives). This data is obtained from a set of ICUs including medical, surgical, coronary care, and neonatal in a single tertiary teaching hospital in the 2001 to 2008 time period. It includes patient information that falls into various categories such as general, physiological, medications, fluid balance, notes, and reports, see Table 1. The entire database totals to about $\approx 27GB$ and contains information about tens of thousands of ICU patients. In order to visualize MIMIC II clinical data within our framework, the relevant patient hierarchies and meta-graph structures need to be computed and stored. This is done in an offline preprocessing step.

First, for each day in a patient's hospital stay, patient details available in the database (e.g., admission ID, age, gender, race, ICD-9 code, drug code, hospital stay length, mean heart rate, mean temperature, and max urine output) are extracted, which results in 38291 patient admissions from 32536 patients. These details are then aligned to make sure all admissions fall on the first time step of the resultant data set. The resulting data set after aligning con-

tains 174 time steps (i.e., 174 days).

Patients in each time step are then clustered together using the metric of [18]. This patient similarity metric was previously applied to the same MIMIC II clinical database to identify patient similarities within the first day of the ICU stay [18]. In order to apply the metric to our research, the required clinical, administrative, and categorical variables are extracted from the database for each day within a patient's hospital stay. Next, correlations across patient groups are computed by tracking individual patients within the groups.

Once the patient groups and their correlation details are stored in our data format, the total data size is reduced to $\approx 680MB$. By precomputing the patient hierarchies and meta-graph structures and storing them using optimized data structures, we allow interactive exploration of patient progression over time for several gigabytes of data.

Researchers are provided with the flexibility to vary the patient similarity thresholds and explore the entire parameter space interactively. Such interaction provides an understanding of how patients group together for varying similarity thresholds within a particular day in their hospital stay. Figure 6 shows several examples of patient groups and their progression for 30, 34, 35, 36, and 40 similarity thresholds. As the similarity thresholds decrease, more patients are grouped together, reducing the complexity of the tracking graph. For a specific similarity metric, exploring the full range of similarities enables researchers to gain insights on that metric's range of values. In our case, upon exploration we realized that for this particular patient similarity metric, the appropriate similarity threshold range is 30-38. Any similarity threshold below or above that range either grouped all patients into one group or divided each patient to be in a separate group.

Our system presents a global concise view of patient progression over time using tracking graphs. The full tracking graph showing the patient progression over time at 36 similarity threshold is displayed in Figure 7. By observing these tracking graphs, specifically feature track length indicating the hospital length of stay of patients, it is clear that although many of the patient stays are less than 90 days (i.e., 3 months), our data set also contains several longer patient stays. Of 32536 patients, we found 6 patients with hospital lengths of stay greater than 90 days.

More importantly, various simplification options available in our system allow researchers to further simplify the tracking graphs. For example, filtering the tracking graph by correlation amount allows removal of the least frequent patient progression paths from the tracking graph, making frequent patterns more prominent. If an analysis is to be conducted only on longer hospital stays of patients, filtering options available within our system, specifically filtering tracking graphs by the length of a feature track, are useful. Figure 8 illustrates several such simplification results.

Additionally, the patient view of our system is useful for obtaining an overview of patient groups. As the user selects a certain patient group, this view displays the details of its patients. The word cloud visualization provides a quick visual overview of the information within a selected patient group, see Figure 4(d). The numerical attributes such as age and hospital length of stay are converted to ranges to obtain more intuitive results. The exact patient details are also presented in the textual visualization within the system, see Figure 4(e).

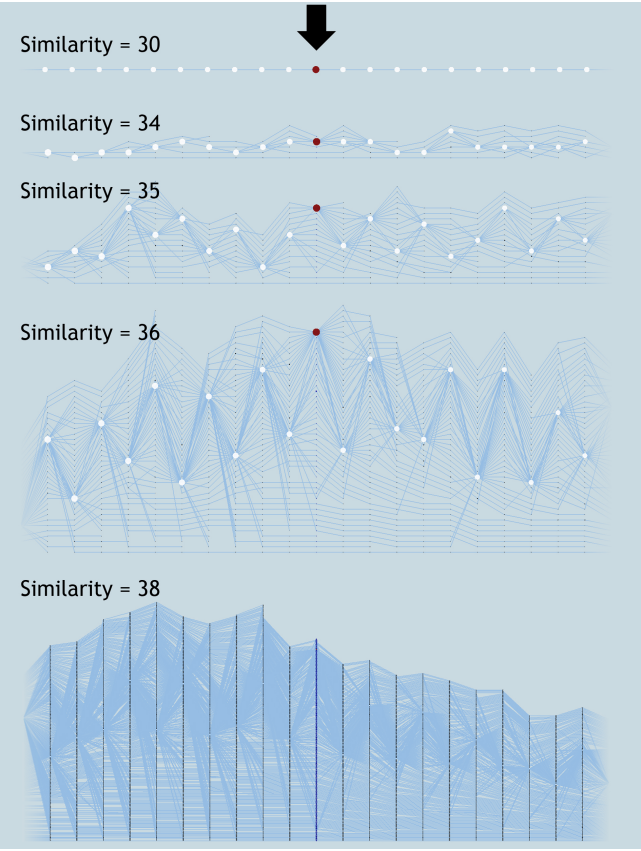

Figure 6: Effects of varying the similarity threshold to explore the temporal progression of patients. Here, patient groups and a portion of their corresponding tracking graphs are shown at 30, 34, 35, 36, 38 similarity thresholds. The focus time step of the tracking graphs is indicated with a black arrow, and the nodes are scaled based on the patient group's size. In each graph, patient progression for 10 time steps both forward and backward in time from the focus time step is displayed.

5 Conclusion and Future Work

In this work, we present a visualization and analysis environment for understanding patient progression over time. The system's interactive abilities to explore patient progression for different similarity metrics and for varying similarities are a distinct advantage over existing techniques used in healthcare. Using our system, researchers are able to explore how patients group together and progress over time, identify frequent progression paths, and also refer back to the native space of data for a visual understanding. By combining optimized data structures and progressive visualization techniques, we enable interactive exploration of terabytes size data, which provides the platform to use this type of analysis in a hospital setting.

Within this work, an existing patient similarity metric is utilized for defining patient similarities. At each moment in time, patient similarities are computed by looking at a patient's current clinical, administrative, and categorical information. A better similarity metric would be one that considers both the current information of the patient and the entire history starting from the hospital admission time. In order to obtain better results, we hope to utilize such a similarity metric in the future. In this work, we demonstrate the applicability of our approach using a publicly available ICU database. We are looking into obtaining additional healthcare databases to use within our system, specifically, databases with geospatial information for which the patient view within our system would prove to be more beneficial. Finally, we aspire to use our visualization and analysis environment in a real-time setting to assist the decision-making process of our collaborating physicians and clinical researchers.

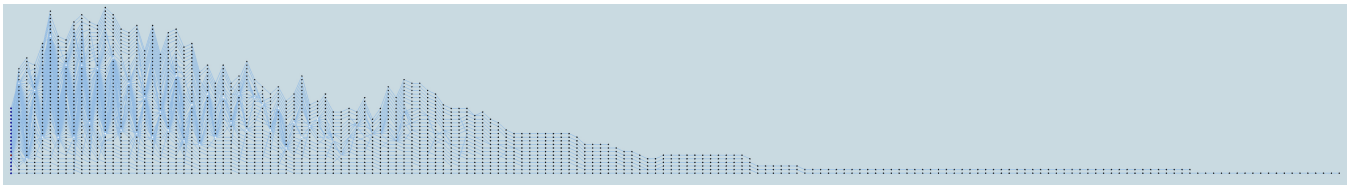

Figure 7: The entire tracking graph showing the complete patient progression for the MIMIC II clinical database. The graph contains 1110 nodes and 1288 edges for a total of 174 time steps. Here, 36 similarity threshold is used.

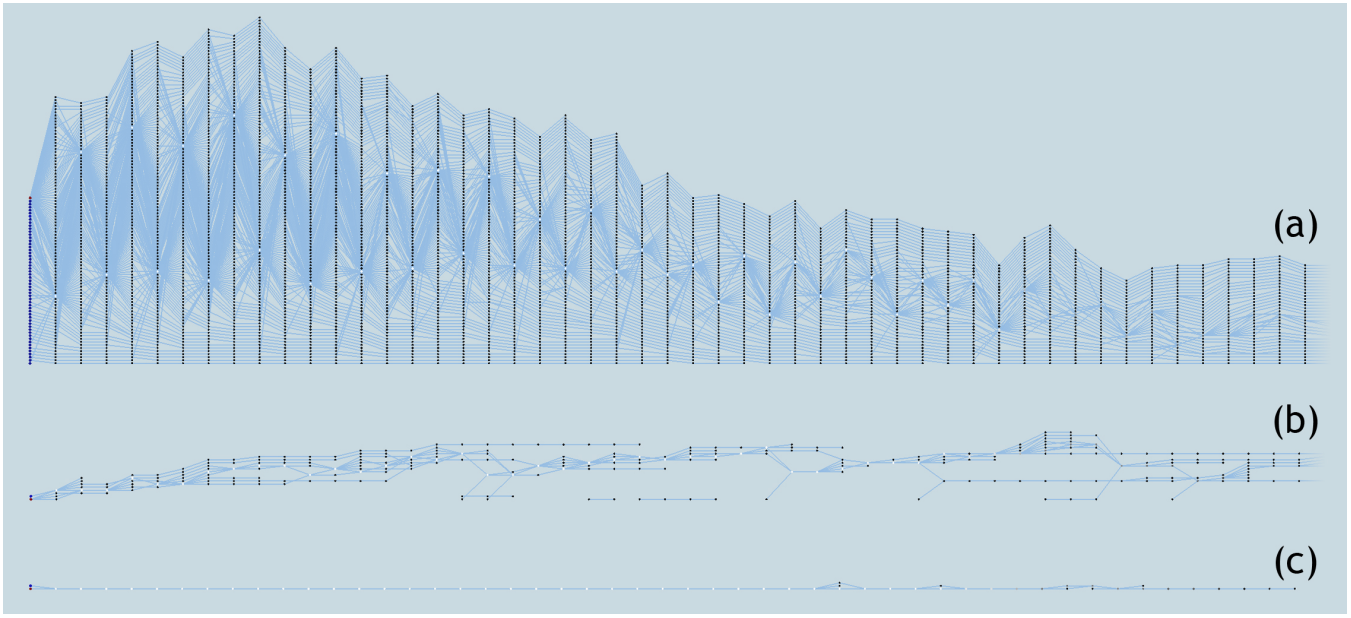

Figure 8: Simplifications of tracking graph. (a) A tracking graph showing the patient progression for the first 50 days within the hospital stay at 37 similarity threshold. (b) Tracking graph in (a) filtered to contain only correlations with overlap$\geq$2. (c) Tracking graph in (a) filtered to contain only patient groups with size$\geq$5.

Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program. This work was also performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-CONF-703800). This work also received funding support from the National Science Foundation Industry & University Cooperative Research Program (I/UCRC) # IIP-1439668.

References

[1] L. Anthony Celi, R. G. Mark, D. J. Stone, and R. A. Montgomery. "big data" in the intensive care unit. closing the data loop. *American journal of respiratory and critical care medicine*, 187(11):1157–1160, 2013.

[2] Awesomium Team. *Awesomium : HTML UI Engine*, 2013. *Available at* `http://www.awesomium.com/`.

[3] S. Bandyopadhyay and E. J. Coyle. An energy efficient hierarchical clustering algorithm for wireless sensor networks. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 3, pages 1713–1723. IEEE, 2003.

[4] P.-T. Bremer, G. Weber, V. Pascucci, M. Day, and J. Bell. Analyzing and tracking burning structures in lean premixed hydrogen flames. *IEEE Trans. Vis. Comp. Graph.*, 16(2):248–260, Mar. 2010.

[5] P.-T. Bremer, G. Weber, J. Tierny, V. Pascucci, M. Day, and J. Bell. Interactive exploration and analysis of large-scale simulations using topology-based data segmentation. *IEEE Trans. Vis. Comp. Graph.*, 17(9):1307–1324, 2011.

[6] D. Cai, X. He, Z. Li, W.-Y. Ma, and J.-R. Wen. Hierarchical clustering of www image search results using visual, textual and link information. In *Proceedings of the 12th annual ACM international conference on Multimedia*, pages 952–959. ACM, 2004.

[7] H. Carr. *Topological manipulation of isosurfaces*. PhD thesis, Citeseer, 2004.

[8] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(5):603–619, May 2002.

[9] J. Davies. *d3-cloud : Word Cloud Layout*, 2012. *Available at* `https://github.com/jasondavies/d3-cloud/`.

[10] W. H. Day and H. Edelsbrunner. Efficient algorithms for agglomerative hierarchical clustering methods. *Journal of classification*, 1(1):7–24, 1984.

[11] J. Ellson, E. Gansner, L. Koutsofios, S. North, G. Woodhull, S. Description, and L. Technologies. Graphviz — open source graph drawing tools. In *Lecture Notes in Computer Science*, pages 483–484. Springer-Verlag, 2001.

[12] Google Team. *Google Maps*. Google, 2005. *Available at* `http://www.google.com/maps/`.

[13] W. Horn, C. Popow, and L. Unterasinger. Support for fast comprehension of icu data: Visualization using metaphor graphics. *Methods Archive*, 40(5):421–424, 2001.

[14] W.-H. Hsu, J. Mei, C. D. Correa, and K.-L. Ma. Depicting time evolv-

ing flow with illustrative visualization techniques. In *ArtsIT*, pages 136–147, 2009.

[15] M. Iwayama and T. Tokunaga. Hierarchical bayesian clustering for automatic text classification. In *Proceedings of the 14th international joint conference on Artificial intelligence-Volume 2*, pages 1322–1327. Morgan Kaufmann Publishers Inc., 1995.

[16] A. Joshi, J. Caban, P. Rheingans, and L. Sparling. Case study on visualizing hurricanes using illustration-inspired techniques. *IEEE Trans. Vis. Comp. Graph.*, 15(5):709–718, 2009.

[17] A. Joshi and P. Rheingans. Illustration-inspired techniques for visualizing time-varying data. In *Proceedings of the conference on Visualization '05*, VIS '05, pages 679 – 686. IEEE Computer Society, oct. 2005.

[18] J. Lee, D. M. Maslove, and J. A. Dubin. Personalized mortality prediction driven by electronic medical data and a patient similarity metric. *PloS one*, 10(5):e0127428, 2015.

[19] M. Monroe, K. Wongsuphasawat, C. Plaisant, B. Shneiderman, J. Millstein, and S. Gold. Exploring point and interval event patterns: Display methods and interactive visual query. *University of Maryland Technical Report*, 2012.

[20] C. Muelder and K.-L. Ma. Rapid feature extraction and tracking through region morphing. Technical report, Citeseer, 2007.

[21] F. Murtagh. A survey of recent advances in hierarchical clustering algorithms. *The Computer Journal*, 26(4):354–359, 1983.

[22] C. F. Olson. Parallel algorithms for hierarchical clustering. *Parallel computing*, 21(8):1313–1325, 1995.

[23] V. Pascucci, P.-T. Bremer, A. Gyulassy, G. Scorzelli, C. Christensen, B. Summa, and S. Kumar. *Advances in Parallel Computing, Cloud Computing and Big Data*, volume 23, chapter Scalable Visualization and Interactive Analysis Using Massive Data Streams, pages 212–230. IOS Press, 2013.

[24] V. Pascucci, G. Scorzelli, P.-T. Bremer, and A. Mascarenhas. Robust on-line computation of reeb graphs: simplicity and speed. In *ACM Transactions on Graphics (TOG)*, volume 26, page 58. ACM, 2007.

[25] V. Pascucci, G. Scorzelli, B. Summa, P.-T. Bremer, A. Gyulassy, C. Christensen, S. Philip, and S. Kumar. *High Performance Visualization: Enabling Extreme-Scale Scientific Insight*, chapter The ViSUS Visualization Framework. Chapman & Hall/Crc Computational Science, 2012.

[26] 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*, volume 2007, page 598. American Medical Informatics Association, 2007.

[27] 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*, page 76. American Medical Informatics Association, 1998.

[28] F. Reinders, F. H. Post, and H. J. W. Spoelder. Visualization of time-dependent data using feature tracking and event detection. *The Visual Computer*, 17:55–71, 2001.

[29] A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman. Interactive information visualization to explore and query electronic health records. *Foundations and Trends in Human-Computer Interaction*, 5(3):207–298, 2011.

[30] M. Saeed, M. Villarroel, A. T. Reisner, G. Clifford, L.-W. Lehman, G. Moody, T. Heldt, T. H. Kyaw, B. Moody, and R. G. Mark. Multiparameter intelligent monitoring in intensive care ii (mimic-ii): a public-access intensive care unit database. *Critical care medicine*, 39(5):952, 2011.

[31] R. Samtaney, D. Silver, N. Zabusky, and J. Cao. Visualizing features and tracking their evolution. *Computer*, 27(7):20–27, July 1994.

[32] B.-S. Sohn and C. Bajaj. Time-varying contour topology. *IEEE Trans. Vis. Comp. Graph.*, 12(1):14–25, 2006.

[33] 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*, pages 457–466. ACM, 2008.

[34] S. D. White and C. S. Frenk. Galaxy formation through hierarchical clustering. *The Astrophysical Journal*, 379:52–79, 1991.

[35] W. Widanagamaachchi, C. Christensen, P.-T. Bremer, and V. Pascucci. Interactive exploration of large-scale time-varying data using dynamic tracking graphs. In *IEEE Symposium on Large Data Analysis and Visualization (LDAV), 2012*, pages 9–17, 2012.

[36] W. Widanagamaachchi, P. Klacansky, H. Kolla, A. Bhagatwala, J. Chen, V. Pascucci, and P.-T. Bremer. Tracking features in embedded surfaces: Understanding extinction in turbulent combustion. *IEEE 4th Symposium on Large Data Analysis and Visualization (LDAV)*, 2015.

[37] K. Wongsuphasawat and D. Gotz. Outflow: Visualizing patient flow by symptoms and outcome. In *IEEE VisWeek Workshop on Visual Analytics in Healthcare, Providence, Rhode Island, USA*, pages 25–28. American Medical Informatics Association, 2011.

[38] K. Wongsuphasawat, J. A. Guerra Gómez, C. Plaisant, T. D. Wang, M. Taieb-Maimon, and B. Shneiderman. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1747–1756. ACM, 2011.

[39] Y. Zhao and G. Karypis. Evaluation of hierarchical clustering algorithms for document datasets. In *Proceedings of the eleventh international conference on Information and knowledge management*, pages 515–524. ACM, 2002.

Understanding Care Plans of Community Acquired Pneumonia Based on Sankey Diagram

Shunan Guo[1,2], Ph.D Student, Chaoguang Lin[2], Bachelor,
David Gotz[3], Ph.D, Bo Jin[1], Ph.D, Hongyuan Zha[1], Ph.D
Linhua Shu[4], M.D, Nan Cao[2], Ph.D
[1]ECNU, ShangHai, China; [2]NYU, ShangHai, China;
[3]UNC, NC, USA ; [4]SHChildren, ShangHai, China

Abstract

A care plan is a sequence of medical interventions formulated for curing a specified disease. Doctors usually craft different care plans for different patients based on their knowledge and experience as well as following the clinical guidelines. An intuitive summarization of previous successful care plans not only provides doctors a reference of successful treatment guidelines, but also helps with the detection of anomalous treatments and the improvement of existing care plans. However, producing such kind of summarization is challenging due to the complexity of the data, i.e., large, temporal oriented, and multidimensional. In this paper, we propose a Sankey diagram-based visualization design to visually summarize care plans based on our medical collaborators' requirements. We apply our tool to a medical dataset of pneumonia patients collected from a Children's hospital in Shanghai, China. Based on the visualization results, doctors detected many interesting findings, which will be discussed in the paper.

1. Introduction

Medical care planning is a critical step in the care delivery process, contributing directly to a patient's treatment pathway and associated outcomes. A care plan is a sequence of medical interventions formulated to help patients manage their diseases. It is often crafted by doctors after a patient is diagnosed according to their intuition and experience, and is often based on recommendations in clinical guidelines. However, the correct care plan is not identical for everyone with a given medical condition due to the individual differences between patients. Thus, in order to help patients obtain better treatment, it is often necessary for doctors to customize the treatment plans recommended in clinical guidelines.

The care plan customization process can be aided, in part, by examining the efficacy of alternative care pathways as experienced by previously treated patients. However, there are several complex issues which make this form of care plan analysis hard to accomplish. First, the large numbers of patients as well as the vast variety in drug types and other variables in medical data increase the complexity of datasets. These issues make the comparison and aggregation over large numbers of patients difficult if not impossible with traditional tools. Second, care pathway data contains temporal events, such as all medications and procedures, which make such data even more difficult to analyze. Third, it is important to not only consider the pairwise correlation of variables across patients, but also the combination of these variables (e.g., multiple medications or diagnoses, often overlapping in time).

Although there are some existing solutions [6-8] for displaying time-oriented medical data, many fail to handle the case of simultaneous or overlapping medical treatment. The typical serial time-oriented data visualization techniques often used for medical data are therefore not quite suitable for such situations. There are also methods that have focused specifically on representing medical care plans [1-4], which inspired our work. However, unlike our work, most of these methods formulate care plans based on existing treatment guideline instead of medicines, thus missing details such as when did what medicine has been used for curing a disease.

In this paper, we present a visual encoding schema based on the Sankey diagram to enable the analysis of treatments for a cohort of patients to examine differences in outcomes for variations in care plans. In this visualization, we summarize a patient's care plan based on a layered graph model in which each node indicates a medicine (i.e., a treatment) and different graph layers indicates different time and links are used to connect the same treatment at the different time. Similar path plans of different patients are aggregated together to form patient cohorts that facilitate the comparison between pathways for finding better or worse medical care plan in a single clinical stage or the entire course of disease. Rich interactions are also designed to support an efficient data exploration and filtering. Throughout the paper we describe our design using a motivated problem related to Community Acquired Pneumonia (CAP). We include five sample analysis after discussed with CAP experts to explain the insights found through the visualization. The major contributions of this paper are as follows:

1. Apply Sankey Diagram to real-world medical care plan data and verified its applicability after the evaluation of domain experts.
2. A novel visual encoding method based on Sankey Diagram to display electronic medical care plan records.

The rest of the paper is organized as follows. We review the related work in section 2 and describe our motivation problem in section 3. We introduce our visual encoding design in section 4 and demonstrate preliminary analysis sample and evaluation results from domain experts in section 5. And finally conclude in section 6.

2. Related Work

In this section we review the papers related to work, which basically includes techniques developed for visualizing temporal event sequences and medical care planning.

2.1 Visualizing temporal event sequences

Generally, this topic lies in the direction of visualizing time-oriented data, which is comprehensively discussed in [5]. We focus on the techniques developed for representing event sequences in the healthcare domain. Sankey Diagram is one of the most intuitive and commonly used methods for representing the event sequence, which is also adopted in our visualization design. Besides this approach, many designs such as LifeLines[6], Timelines[7], Eventflow[8] and many other timeline-based representations[9-12] align event sequence horizontally along a timeline in which one patient record is split into different event categories, thus making them inefficient for cohort analysis. LifeFlow[13] aggregates health records of multiple individuals based on a Treemap but each event type are divided into several pieces and is hard to analyze in general. There are also visualization tools like[14-20] aggregates multiple records based on data transformation and mining techniques, DecisionFlow[21] supports an in-depth analysis of a heterogeneous multidimensional event sequence at different phrases via rich interactions and flow based visualization design. Although powerful, none of these techniques are specifically designed for revealing a care plan in the electronic health records.

2.2 Visualizing medical care planning

Among many visualization techniques developed for representing electronic health records [22-25], visually representing care plans attracts more and more research interests in recent years due to its usefulness for supporting clinic process. Most projects dealing with representing care planning are based on flow-chart algorithms, which is widely known by physicians and requires minimal learning efforts. For instance, CareVis[4] and AsbruView[3] gives visualization solution for large and complex flowchart. GapFlow[1] shows the derivation during different medical treatment plans. Most of these techniques take the existing care plans as the input, which are not always available. Our work is largely inspired by CareFlow[2], which is designed to assist doctors in finding better care plans based on the treatment records. We adopt its visual design to help analyze a group of patients. Different from CareFlow, which shows cohorts' physical outcomes after applying a sequence of treatments, our work improves this design by strictly aligning the treatments along a timeline, which help illustrate the combination or correlations among different medicines used at the same and different time.

3. Motivation Problem and Dataset

Community-acquired pneumonia (CAP) is one of the most common infectious diseases and has been recognized as a potentially lethal condition for nearly two centuries[26]. It is also a serious infection that afflicts children throughout the world. The average annual incidence of pneumonia in children younger than 5 years of age is 34-40 cases per 1000, and is increasing every year, generally becoming the largest killer of children[27]. Symptoms suggestive of pneumonia basically include 80 percent of fever combined with respiratory symptoms such as cough, sputum production, pleurisy, and dyspnea. The pathogens responsible for community-acquired pneumonia in children mainly includes mycoplasma, influenza virus, and bacteria. However, there have been few attempts to devise treatment guidelines in China. Guidelines from North America and Europe are not practically useful enough for doctors in China due to different infection environment and etiologic process. And the treatment guideline for CAP in China is lack of support of population statistics, or slightly out-of-date because of pathogenic variation. Besides, the antibiotic abuse is widespread is the medical treatment of CAP in China, which results in greater potential hazards.

With the above issues in mind, our medical collaborators are interested in analyzing the past clinical cases, in order to find care pathways that are mostly used or those anomalous or irrational ones, and their corresponding outcomes. Furthermore, doctors also wish to find the proportion of patients for different pathogens and their correlated physical signs, as well as physical reactions to different drugs. To address these problems, we are given access to a dataset of

nearly 3,000 children patients with pneumonia and with various medical treatments based on nearly 100 types of medications. All these patients were hospitalized. The patients' anonymized information, prognosis results (i.e., pathogen types), the full treatment history (e.g., medicines) during the hospitalization, and the monitoring of the patients' body temperatures was also given.

4. Visualization Design

In this section, we first present certain design tasks stemmed from our discussion with two doctors. We then provide a detailed description of our data aggregation and visual encoding methodology. Finally, we introduce a couple of user interactions for data exploration through graph manipulation.

4.1 Design Tasks

Clinical CAP physicians are often faced with the difficulty of making precise care plans for different patients. Currently, an efficient care plan can only be made by these doctors manually by using their own experiences and domain knowledge to bridge the outcomes to the corresponding medical treatments. This procedure is usually extremely inefficient and time consuming. Therefore, the doctors expected a tool or even a system which could provide references from the existing and similar care plans to support their decision making process. This is especially important in China as most of the doctors need to handle hundreds of patients every day. The doctors are particular need a system that can help them automatically integrate the treatment with the outcomes and show existing care plans that are related to a focal disease for their reference. They hope this tool can also help to summarize the existing care plans so that they can easily identify which one is more efficient in terms of curing the disease. To meet their requirements, we compile a list of visualization design tasks as follows.

T1 Integrating the medical treatments and outcomes of each patient. Both medical treatments and the outcomes such as the monitoring results of the body temperatures are recorded independently over time. Therefore, for each patient, we need to align different data records based on their timestamps to build a medical care path so as to identify the order of medicine being used and corresponding physical sign afterwards.

T2 Clustering patients with similar medical care paths. In order to proceed care plan comparison and cohort analysis, our visualization needs to gather similar paths and show average physical outcomes. In this case, different care plans are automatically distinguished, saving the doctors from looking through thousands of individual pathways.

T3 Revealing the statistical details of each treatment. The statistical details of treatment include the number of involved patients and the frequency of a medicine being taken. This helps doctors to quickly differentiate the critical care plan for groups of patients and anomalous plans for different patients.

T4 Associating care plans with pathogens. Finding correlations between the treatments and different pathogens are considered to be very important as it will help doctors to make a correct decision at the early stage. Therefore, the proposed system should be able to differentiate different pathogens of a focal disease and associate it with different care plans. More precisely, the system should be able to illustrate the proportions of patients for each pathogen type and their corresponding reaction after taking each treatment.

T5 Facilitating visual data filtration. Considering the actual care pathways can be chaotic due to the complicated situations the doctors may face, the system should enable doctors to filter exceptional cases and explore data through interactions in order to reveal the main patterns.

T6 Easy browsing of raw data. The raw data, such as the name of a medicine, a pathogen type, and specific amount of the cohort's size can help doctors to a better understanding a care plan. Some other features of the patient, such as the lab test results, the exact date of hospital admission and discharge can be of great value for doctors to perceive the essence beyond the visual representations. Thus the visualization should enable analysts to explore raw data easily.

4.2 Data Model

With the above dataset and design tasks, our goal is to find and differentiate care plans that are used among groups of patients and compare their outcomes. We approach this goal by aggregating all the patient records based on how the treatments are performed. Specifically, patients take the same medicines at the same stage (defined by the number of days in hospital) are grouped together. Thus the grouping results at different stages forming a summary of the treatment history as shown in Figure 1. In this data model, nodes indicate medicines and links indicate groups of patients, which have two primary

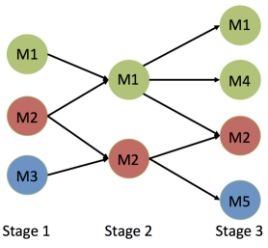

Figure 1.The data model for summarizing the care plans

attributes: the size of the patient group and the outcome (e.g., body temperature) after taking the previous treatment.

As a result, the final directed-acyclic-graph captures every existing medical care pathway as well as the number of patients flowing through the same sub-paths, and the average body temperature during each time interval. This graph serves as the prototype and preparation of our Sankey Diagram implementation.

4.3 Visual Encoding

The above data model can be intuitively visualized by a Sankey diagram and the aforementioned tasks guide our design of the visual encoding schema. Figure 2 illustrates our visual encoding design. It illustrates (a) the proportion of patients taking each medicine at the same time interval, (b) the proportion of patients for each pathogen type, and (c) the average body temperature for each time interval after or before a treatment was taken. In particular, this visualization consists of several components, which are described as follows:

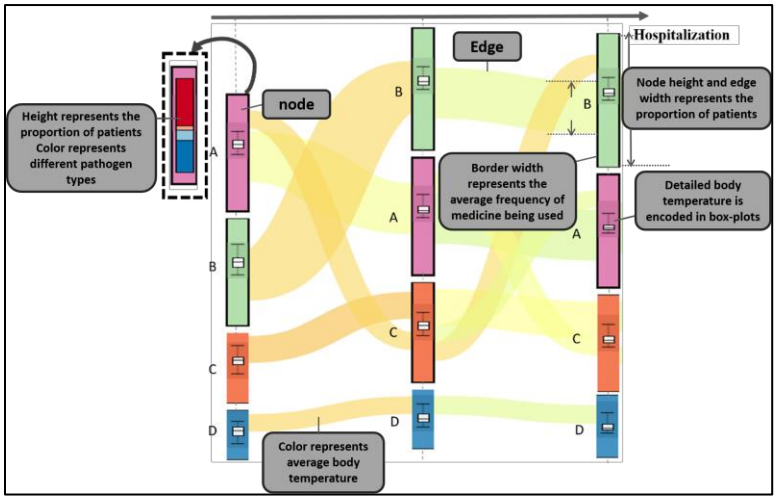

Figure 2.Nodes represent treatments and are positioned horizontally in temporal order. Edges connect treatments in each day and are color-coded to represent the average body temperature.

Nodes (Treatment): Nodes represent treatments and are positioned along the horizontal axis indicating treatment sequence over time. With different filling color representing different kind of medicine and the height representing the proportion of patients taking a given medicine (**T3**). And the border width of each node represents the frequency of this medicine being taken in a day. For example, as illustrated in Figure 2, there is a total of four types of medicine being taken during hospitalization, A, B, C and D. Medicine A and B are used by large and almost equal proportion of patients, while medicine D is used by the smallest cohort.

Layer: The overall view is horizontally divided into several layers according to the maximum days of hospitalization. Each layer represents one day of hospitalization, showing the treatment information from the first day in the hospital until the last. And the nodes of each layer represent the treatments being taken on that day. For example, as shown in Figure 2, the longest hospitalization of this group of patients is three. And for each day, there are all four types of medicine being taken. The vertical adjacency relation of each pair of nodes represents the correlation between medicines, which was implemented using force-directed layout with back and forward propagations. In other words, medicines that are frequently used as combination are more likely to be shown in the neighborhood. For example, the position of medicine A and B switches vertically in day 2, probably because the amount of people using medicine A and C slightly increases edging out the combination of medicine B and C.

Edge (Temperature): The edges are colored according to the average body temperature of all patients represented by the corresponding source node. Elements that are colored red represents parts of the care plan where patients are in the state of fever, whereas elements colored green are care plans where the patients' body temperature becomes normal. The width of edges is correlated with the source nodes and the target nodes, representing the number of patient flows through this particular edge.

Node Encoding: There are two types of filling for each node and doctors can switch one to another through interaction. The first is box-plot, which shows the variance of body temperature in the group of patients and to some extent avoids the temperature from being balanced through averaging. The box-plot is displayed in a shadowed area with certain height, and is visible only when the cohort size reach the threshold in order to maintain its comparability. The other filling is the pathogen proportion. Each color represents a type of pathogen and the height is proportional to the number of patients (**T4**).

4.4 User Interactions

As expressed in our design tasks (**T5**), user interactions are key to data and pattern exploration. Our visualization allows doctors to interactively perform the following actions.

Highlighting: As shown in Figure 3, doctors can highlight the overlap edges of the corresponding group of patients through hovering mouse over an edge, a node or a type of pathogen. This helps doctors to ascertain the actual path that the cohort goes through.

Hovering: Hovering also triggers the display of tooltips which provides the exact name of a type of medicine or pathogen as well as the specific number of patients (Figure 3).

Filtering: Doctors can filter data in two ways. First, doctors can select a cohort of patients and construct a new Sankey Diagram through double clicking on edges and nodes. This can help doctors to get a better analysis on the particular cohort they are interested in. Second, doctors can use the bidirectional slider to filter both nodes and edges on the number of patients to remove small subgroups. The graph will then show only the care plans that are mostly taken by removing thin edges and unconcerned nodes. Also the doctors can filter thick edges to observe the exceptional cases.

Raw data exploration: Doctors can filter the items displayed in raw data list through the filtering box on the top, or by clicking an edge or a node to select a cohort of patients.

Encoding Switching: As described in 4.3, doctors can choose whether to display pathogen information or not. And the fillings of the nodes can be changed through the switch on top.

5. Analysis of Pneumonia Care Plans

5.1 Domain Expert Interview

Based on the above visualization design, we demonstrate our system in front of a medical expert team based on a core dataset cleaned from the raw data. The medical expert team, led by the director of the department of respiratory disease, has rich experiences in both CAP clinical and treatments. The core dataset used for the system evaluation and expert interview consists of 953 patient records with a complete using records of 24 types of medicines collected in one year. The interview starts with a tutorial covering the visualization design and interactions. We then ask doctors to use the system on their own for exploring the core dataset. After a full understanding of the system's capability, each doctor was given a list of questions as the guideline of system evaluation. Doctors are asked to provide their feedbacks and suggestions or raise any questions during the process of using the system. The interview lasted approximately 2 hours. We recorded the entire conversation, and took notes of their comments. The rest of the section will give a brief summary of doctors' comments. And we list some of the medical related findings in section 5.2 and doctors' suggestions on system improvement in section 5.3.

Both doctors are very impressed by the volume of information that the visualization provided. They commented that "problems can be displayed intuitively through the view", and "Comparing to traditional statistical analysis method, this aggregates different types of statistics data with more detailed value in each hospitalization day". Even though they mentioned that "since the system provides such huge amount of information, it will take some learning efforts to master the use of the system". However, they believe "this tool can be very convenient and efficient in exploring large datasets once you get used to it". In addition, the first doctor suggested that "we should popularize this system to a larger platform in response to the trend of 'Precision Medicine'". He also expressed his alacrity of providing more data for comparison to see whether the patterns of care plan evolves through years. The second doctor is particularly fond of the idea of displaying care plans and all sorts of statistical data in this way. She said, "We used to display statistical data in forms of pie charts or histograms. However, this system is much more powerful. It is not only capable of displaying multiple types of statistical data integrally, but even the way they change over time."

5.2 Preliminary Analysis Result

In this section, we describe the doctors' feedback and demonstrate several interesting findings detected by the doctors.

F1 The system is effective in displaying major care plans. As shown in Figure 3, when the doctor hovers on the edge from medicine A (t1) to medicine A(t2), edges from and to medicine B are also highlighted. Considering the width of highlighted edges, they found that nearly 50% of the patients taking medicine A at day 2 and 3 are also taking medicine B at the same time. In particular, medicine A is Ceftriaxone and medicine B is Azithromycin. The combination of Cephalosporin drugs and Azithromycin is a common approach for dealing with children respiratory

diseases and works pretty well. The same pattern is also found with medicine C, which is another type of Cephalosporin drug.

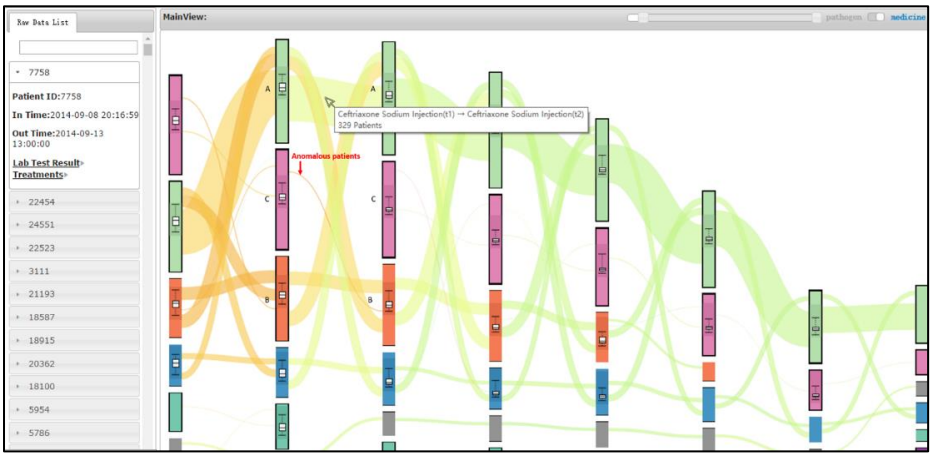

Figure 3. Overlap edges are highlighted after hovering, showing how the care pathways flow through the rest of the graph. Edges between node A and B are also highlighted after hovering the edge from A to A, indicating that medicine B are mostly taken along with medicine A. And few edges from medicine C is highlighted indicating medicine A and C are usually not taken together. Doctors can refer detailed information of the anomalous patient through raw data list on the left by clicking the corresponding edge.

F2 Anomalous care plans. The doctors also found several suspicious care plans. For example, as shown in Figure 3, medicine B is seldom used. Sometimes a thin edge comes from medicine B, according to our highlight mechanism, the doctors found a small group of patients take both medicine A and B at the same time. However, both A and B are Cephalosporin drugs, the doctors believed that the mixed used of these medicines are problematic, which worth a further inspection of the raw data.

F3 Verifying statistical results. The doctors are very familiar with the dataset and have already conducted some simple statistical analysis and illustrate the results base on pie charts and histogram. The result showed that 31% of the patients are taking Cefuroxime (medicine C) while 37.15% of the patients are taking Ceftriaxone (medicine A). From the Figure 3, we can see that medicine A and C are used by a large proportion of patients and their heights are substantially equal. This fully coincides with the previous statistics. However, our visualization design also shows the variety of basic statistics over time, which was preferred by the doctors.

F4 Revealing the correlation between pathogen and treatments. As illustrated in Figure 4, the doctors switched fillings of the nodes into pathogens proportions and they found a large proportion of mycoplasma patients taking Cephalosporin drugs. This is a surprising finding as Cephalosporin drugs take no effect in killing mycoplasma viruses. This also suggested problematic care plans as many doctors didn't take pathogens into consideration while making a prescriptions and most mycoplasma patients are not taking the right medical care plans at the very beginning. Thus, it is not hard to explain why the portion of mycoplasma patients maintain fever until the fourth day of hospitalization, while the patients taking Amoxicillin Sodium and Potassium recover in the second and third days (Figure 5).

F5 Revealing the care path of specific cohorts. The doctors also inspected the treatments used by a small group of patients used Meropenem and had extremely high fever and were cured after a long treatment. As shown in Figure 6, most patients recovered on the 7th day and one patient was still in the fever. In addition, Meropenem is an advanced antibiotic drug, which were only used on patients with very severe symptoms. Thus, this specific cohort must have very serious disease and are not easy to recover.

5.3 Discussion

Apart from patterns observed in the system, doctors also provide their suggestion on improving the system. First, they think it's necessary to provide filter for doctors to select patients with a certain type of pathogen. Because the symptom of pneumonia is very much pathogen-related. Sometimes doctors are more interested in analyzing the features of one specific pathogen in order to improve its corresponding care plans. Second, they believe the pre-selection of medicines is also needed. Since in most cases, doctors are only interested in a particular group of medicine. For example, during the interview, both doctors are very eager to find out the usage of Cephalosporin medicines, however, there are two Cephalosporin medicines and the system does not provide medicine selection. Even though they can capture most of

the information from the general view, they believe patterns will be revealed more clearly if those two medicines were displayed alone.

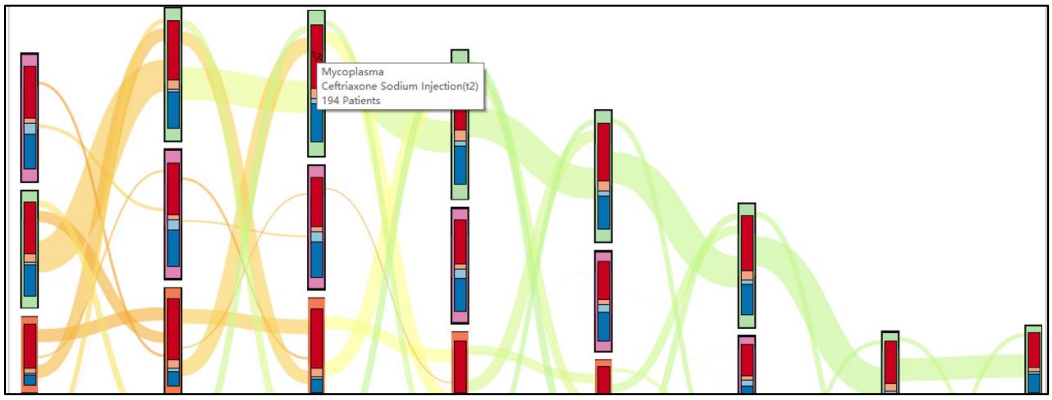

Figure 4. Mycoplasma patients account for a large proportion of each medicine type and usually hospitalize with high body temperature.

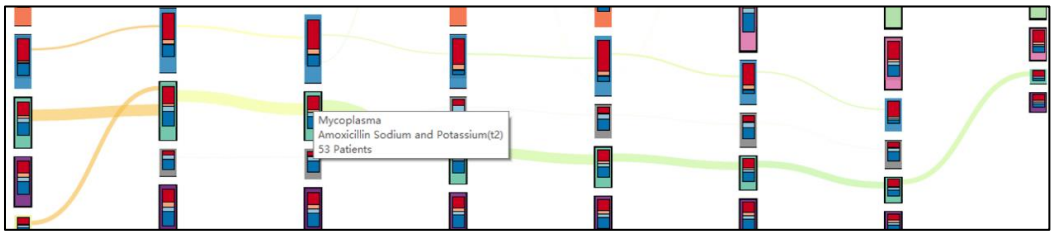

Figure 5. Mycoplasma patients taking Amoxicillin Sodium and Potassium turns out to recover faster than those taking Cefuroxime comparing to figure 4 due to mycoplasma virus's immunity to Cephalosporin drugs.

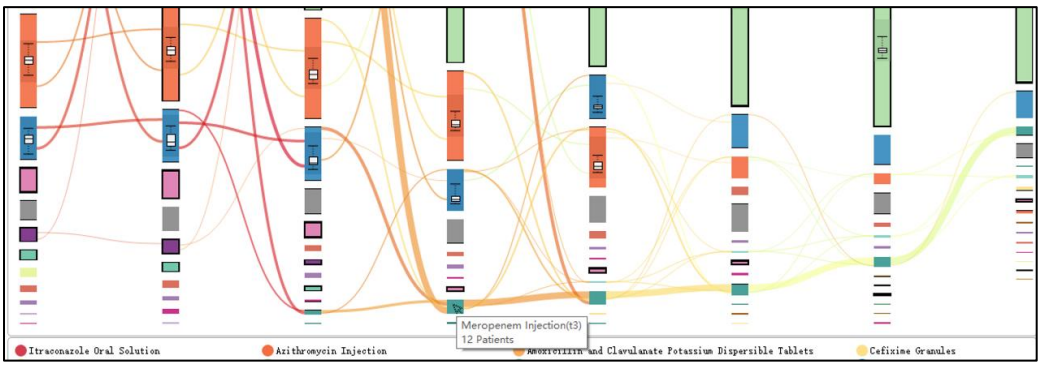

Figure 6. Patients taking Meropenem trends to have extremely high body temperature when hospitalize, and take long time to recover. The observation is confirmed by the fact that Meropenem is an advanced type of antibiotic and is only taken by patients with very severe symptoms.

6. Conclusion

In this paper, we presented techniques for representing medical care plans based on Sankey diagram. In our design, we aggregate different treatments of a focal disease into a layered directed graph. In this graph, each node represents a medicine (i.e., the treatment) and links are used to connect the same medicines used at different time and encode the corresponding outcomes (e.g., body temperatures). Our visualization has been used to summarize care plans of community-acquired pneumonia based on a patient dataset collected from a children's hospital in China. Based on the visualization, many findings such as anomalous care plans were detected by doctors, which verifies the usefulness of the tool. The future work includes proposing advanced algorithm for visual clutter reduction and improving the design to illustrate the situation in which multiple medicines are taken.

Acknowledgements

This work is supported by the Shanghai Committee of Science and Technology, China (Grant Nos. 15JC1401700).

References

1. Gotz, D., et al., GapFlow: Visualizing Gaps in Care for Medical Treatment Plans. IEEE VIS Poster (2013), Atlanta, Georgia.
2. Perer, A. and D. Gotz. Data-driven exploration of care plans for patients. in CHI'13 Extended Abstracts on Human Factors in Computing Systems. 2013. ACM.
3. Kosara, R., et al., AsbruView: capturing complex, time-oriented plans—beyond flow charts, in Diagrammatic Representation and Reasoning. 2002, Springer. p. 535-549.
4. Aigner, W. and S. Miksch, CareVis: integrated visualization of computerized protocols and temporal patient data. Artificial intelligence in medicine, 2006. 37(3): p. 203-218.
5. Aigner, W., et al., Visualization of time-oriented data. 2011: Springer Science & Business Media.
6. Plaisant, C., et al. LifeLines: visualizing personal histories. in Proceedings of the SIGCHI conference on Human factors in computing systems. 1996. ACM.
7. Harrison, B.L., R. Owen, and R.M. Baecker. Timelines: an interactive system for the collection and visualization of temporal data. in Graphics Interface. 1994. Citeseer.
8. Monroe, M., et al., Temporal event sequence simplification. IEEE transactions on visualization and computer graphics, 2013. 19(12): p. 2227-2236.
9. Aigner, W., et al. PlanningLines: novel glyphs for representing temporal uncertainties and their evaluation. in Ninth International Conference on Information Visualisation (IV'05). 2005. IEEE.
10. Bade, R., S. Schlechtweg, and S. Miksch. Connecting time-oriented data and information to a coherent interactive visualization. in Proceedings of the SIGCHI conference on Human factors in computing systems. 2004. ACM.
11. Cousins, S.B. and M.G. Kahn, The visual display of temporal information. Artificial intelligence in medicine, 1991. 3(6): p. 341-357.
12. Karam, G.M. Visualization using timelines. in Proceedings of the 1994 ACM SIGSOFT international symposium on Software testing and analysis. 1994. ACM.
13. Wongsuphasawat, K., et al. LifeFlow: visualizing an overview of event sequences. in Proceedings of the SIGCHI conference on human factors in computing systems. 2011. ACM.
14. Burch, M., F. Beck, and S. Diehl. Timeline trees: visualizing sequences of transactions in information hierarchies. in Proceedings of the working conference on Advanced visual interfaces. 2008. ACM.
15. Vrotsou, K., K. Ellegard, and M. Cooper. Everyday life discoveries: Mining and visualizing activity patterns in social science diary data. in Information Visualization, 2007. IV'07. 11th International Conference. 2007. IEEE.
16. Fails, J.A., et al. A visual interface for multivariate temporal data: Finding patterns of events across multiple histories. in 2006 IEEE Symposium On Visual Analytics Science And Technology. 2006. IEEE.
17. Vrotsou, K., J. Johansson, and M. Cooper, Activitree: Interactive visual exploration of sequences in event-based data using graph similarity. IEEE Transactions on Visualization and Computer Graphics, 2009. 15(6): p. 945-952.
18. Wongsuphasawat, K., et al., Querying event sequences by exact match or similarity search: Design and empirical evaluation. Interacting with computers, 2012. 24(2): p. 55-68.
19. Wongsuphasawat, K. and B. Shneiderman. Finding comparable temporal categorical records: A similarity measure with an interactive visualization. in Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on. 2009. IEEE.
20. Gotz, D., F. Wang, and A. Perer, A methodology for interactive mining and visual analysis of clinical event patterns using electronic health record data. Journal of biomedical informatics, 2014. 48: p. 148-159.
21. Gotz, D. and H. Stavropoulos, Decisionflow: Visual analytics for high-dimensional temporal event sequence data. IEEE transactions on visualization and computer graphics, 2014. 20(12): p. 1783-1792.
22. Frankovich, J., C.A. Longhurst, and S.M. Sutherland, Evidence-based medicine in the EMR era. N Engl J Med, 2011. 365(19): p. 1758-1759.
23. Ebadollahi, S., et al. Predicting patient's trajectory of physiological data using temporal trends in similar patients: A system for Near-Term prognostics. in AMIA Annu Symp Proc. 2010.
24. Gotz, D. and K. Wongsuphasawat. Interactive intervention analysis. in AMIA annual symposium proceedings. 2012. American Medical Informatics Association.
25. Wang, T.D., et al. Aligning temporal data by sentinel events: discovering patterns in electronic health records. in Proceedings of the SIGCHI conference on Human factors in computing systems. 2008. ACM.
26. Bartlett, J.G. and L.M. Mundy, Community-acquired pneumonia. New England Journal of Medicine, 1995. 333(24): p. 1618-1624.
27. McIntosh, K., Community-acquired pneumonia in children. New England Journal of Medicine, 2002. 346(6): p. 429-437

Visualization for All: The Importance of Creating Data Representations Patients Can Use

Carolyn Petersen, MBI, MS
Mayo Clinic, Rochester, MN, USA

Abstract

Complex medical information and data commonly are presented in a variety of graph, chart, and picture formats. Although these presentations perform adequately for health care professionals with medical training, they may be incomprehensible to patients, particularly those with low numeracy and/or graph literacy. However, the advent of learning health care systems, shared decision making that involving patients, greater emphasis on patient engagement, and rising patient advocacy make it essential that patients receive data representations that they can use to assess health risks and select treatment. This article describes the growing need for patient-friendly data visualizations, reviews what has been reported about patients' ability to use information presented in various formats, and identifies approaches to make data representations that are more useful for patients, their families, and caregivers.

Introduction

Approaches to data visualization historically have been driven by scientists and academics seeking to communicate their work to others with similar backgrounds and training. When the roles of patient and medical practitioner were strictly defined and largely separated, these traditional methods of displaying information were effective and sufficient for the audience accessing the information. In the evolving medical environment, however, there is a need to make health information comprehensible to a broader audience, in particular patients and their family and caregivers.

The need to develop methods that support patient visualization of data is driven by multiple circumstances and trends. Within the emerging learning health care systems patients may increasingly face new and potentially different therapeutic options than they have been offered in the past. Although the systems will drive this movement, patients will be expected to assume greater responsibility for understanding their care options and functioning as a part of a continuous feedback loop. The adoption of new treatment protocols may occur more rapidly than in the past, and patients may find themselves expected to properly follow disease management programs about which they have gained little awareness through consumer publications and social media. Effective patient communication and educational aids that clearly illustrate health risks and outcomes will be critical to the success of learning health systems.

At the same time, patients are becoming more interested in learning about their condition and how to manage it and seek information to help them do so; nearly three-quarters of Americans have searched the Internet for health information[1]. The availability of tools such as MedlinePlus and Google Scholar offers access to medical information that is much more technical than the patient education materials commonly offered by providers, creating a need for patients and their families and caregivers to interpret complex information. Too, greater activism by disease-focused service organizations and patient-driven advocacy groups has raised patient interest in knowing more about their health and making informed decisions about their care.

In an environment of shared decision making, patients not only need to better understand their health care needs but also must be better prepared to follow their providers' guidance before, during, and after treatment. Patients frequently are asked to make decisions about treatments, and the greater the uncertainty about which treatment is best, the more likely they will be asked to choose[2]. Data visualization approaches that are easy to understand by patients and their physicians will increase the likelihood that patients make choices with which they are satisfied in both the short and long term.

Improved quality of care and more desirable outcomes have been a goal of both patients and providers for many years[3], and facilitation of greater patient engagement is a key driver of both quality and outcomes. Furthermore, researchers have an ethical obligation to maximize the value of data given by patients[4]. Sharing the information gleaned from that data in formats readily understood by other patients is one way researchers can ensure that this responsibility is met.

Data Visualization among Patients: The Current State

Much of the complex risk and outcomes data patients need to understand comes from peer-reviewed journals, medical meeting presentations, health system-developed materials, and direct-to-consumer advertising. Journal articles and medical presentations typically include data in such forms as Kaplan-Meier curves, forest plots, funnel plots, violin plots, waterfall plots, spider plots, swimmer plots, heatmaps, and circos plots[5], but these representations often are unfamiliar and prove challenging for patients and caregivers to interpret. Newer formats such as network analysis diagrams and transit map diagrams also may be difficult for lay users to understand. Lacking a clear, meaningful expression of data, patients may struggle to use the information effectively in managing their health, particularly if they lack adequate health literacy[6]. New approaches to data visualization can support greater patient engagement and, as a result, improvements in health outcomes.

The ability of patients to understand health information encompasses multiple skill sets including health literacy, numeracy, and graph literacy (sometimes referred to as graphicacy). Literacy involves the ability to comprehend information presented in text, while numeracy refers to the ability to interpret data presented in numbers. Graph literacy is the ability to understand present information in the form of sketches, photographs, diagrams, maps, plans, charts, graphs, and other nontextual, two-dimensional formats[7]. Graph comprehension involves three skills: information extraction, information interpolation and interpretation, and information extrapolation and analysis[8].

Within the health communications field, researchers have focused primarily on concerns related to health literacy and, to a lesser degree, on health numeracy. Graph literacy has been recognized as a separate skill, but has received relatively little attention. Tools commonly used to assess the readability of health information such as the Flesch Reading Ease Readability formula[9] do not take into account features specific to health writing, such as a higher presence of technical terms not common to everyday English, rendering these tools unsuitable for use in designing and assessing visual data representations. The Centers for Medicare and Medicaid Services highlights this deficiency in its toolkit for developing clear written materials, cautioning developers against relying on readability scores when determining how easily readers can interpret health information[10]. Another health information assessment tool, the Centers for Disease Control and Prevention's Clear Communication Index, has been used successfully to evaluate health information infrastructure such as patient portals[11], but also does not address graph literacy.

Numeracy skill may offer some reflection of patients' ability to understand graphical representations of health information, but it is unclear whether numeracy can reliably predict graph literacy. In a study in which highly literate participants interpreted graphs depicting breast cancer risk, made hypothetical decisions about treatment, and indicated preferences for graph format among line, horizontal bar, vertical bar, and iconic graphs, numeracy predicted graph literacy[12]. In another study in which prostate cancer patients performed similar tasks, numeracy and graph literacy were not highly correlated ($r = 0.37$)[13]. Within the low-literacy subgroup of participants in this investigation, graph literacy was more closely correlated ($r = 0.59$-0.90) with the ability to interpret information in a dashboard format. These studies support earlier work suggesting that less numerate individuals may be less able to interpret graphs[14], but other work comparing representations of likelihood in various numerical formats suggested that graphs may be helpful in conveying risk to individuals with low numeracy[15,16].

Though graphs have sometimes been used as a strategy to circumvent the challenge of communicating to people with low health literacy, graphs are not necessarily intuitive and may not be more easily interpreted[17] or result in a greater understanding of health risk when individuals' understanding of risk is accurate prior to viewing a graph[18]. The format of data presentation, including the use of color and icons, the quantity and placement of text and numbers, the relative size of elements within information, and the quantity of information provided also influence how easily patients and consumers comprehend and use information in health decision making. Presentation formats that reduce the cognitive load on users by limiting the number of visual elements and highlighting the most important details improve users' ability to interpret data regardless of health literacy skill and socioeconomic background[19].

Infographics, also known as pictographs, are an increasingly popular approach to data representation, and their appropriate use can enhance consumer awareness and understanding of health-related concerns[20]. An effective infographic draws attention to comprehendible information and narrows the chasm between medical professionals and consumers/patients[21]. Previous investigators have reported pictographs to be the best format for communicating probabilistic information to patients when patients and their providers make decisions together[22], but because pictographs also have been associated with conveyance of relatively straightforward information, they may fail to convey highly nuanced health information.

Visualizing the Future

The design and accuracy deficiencies of mobile health (mHealth) tools as a class of devices have been described[23], and investigations of individual apps and class of apps also note these problems. A clinical trial evaluating a smartphone app marketed as a tool for consumer sleep monitoring found no correlation between the app's readings and readings simultaneously taken during in-laboratory polysomnography[24]. In a review of 28 of the top 200 free and for-pay apps in the "health and fitness" category of Google Play and iTunes, the usability score averaged 13.5 of 20, indicating that some of the apps offered little benefit to users[25]. Quality testing of apps for prevention of driving after drinking[26], pregnancy prevention[27], and other health-related tasks yields similar results.

Fortunately, skillful data visualization offers the potential to avoid these pitfalls. An approach involving multiple user-centered design methods including focus groups, participatory design sessions, and usability evaluation for mHealth applications offers one, though not the only, way forward[28]. Use of an iterative approach not only facilitates a comprehensive assessment of user needs, but also allows patients to describe in their own words what confuses them and whether subsequent designs make data more accessible. Too, use of an iterative approach supports development of a data visualization form that has undergone usability testing and refinement, thereby increasing the likelihood of adoption by patients as well as clinicians.

Incorporating the findings of communication research into a single or small number of formats for data presentation may seem to be an onerous task. One path to success may involve creating familiar graphical displays that allow patients to put to work existing skills, rather than attempting to design data presentations and interfaces based on what people say they like about existing formats[13]. Graphs tend to be more complex than infographics, drawing upon a broader and more developed set of analytical skills than those needed to comprehend the content of illustrations. Creation of graphs using the design principles on which infographics and other less complex images (e.g., furniture assembly instructions) are based merits further exploration.

Patient-provider communication also plays a critical role with regard to shared decision making. Patients' comprehension and ability to participate in treatment and care management increases when patients and their providers engage together with the information, in particular when the meaning of numbers such as lab results are involved[29]. Effective data representations will make it easier for provides to share complex information with their patients, and a commitment to these important conversations will create an environment in which care can become a partnership with improved outcomes as its result.

The growing interest in methods of data visualization that permit new formats and support interactivity[30,31] provides an opportunity to develop visualization approaches that also facilitate patient education and engagement. For example, a tool developed by Weissgerber and colleagues allows users to build interactive line graphs by manually entering data or uploading .csv files[32]. The resulting graphs may be viewed, saved, and downloaded as users require. This ability to enter data and observe changes may be particularly attractive to users who collect data through wearables, in-home sensors, or other devices.

Within the field of patient education, data visualizations often occur within the broader context of decision aids, which combine text, numerical, and graphical elements to help patients answer a particular question, such as "What treatment is most effective for my condition?" Although some work considers specific elements, such as a study assessing how patients with low literacy interpreted quantitative risk presented in different graph formats[33], most evaluation studies of decision aids tend to investigate the effectiveness of the aids as a whole rather than the efficacy of specific elements within the decision aid. Development of research designs that isolate patients' engagement with and understanding of charts and infographics will help patient education professionals develop data visualizations that more effectively communicate complicated information.

In the online environment, dynamic interfaces that present data in different formats depending on the user's role (e.g., clinicians, patients) offer another opportunity to facilitate patient understanding and medical decision making. In one implementation within an electronic medical record, the addition of a presentation layer within the graphical user interface permitted data display that performed desirably for physicians and patients[34]. Both groups were able to achieve high performance with the interface during usability testing, and the EHR-based implementation reduced operation time compared to other EHRs and paper-based processes. The dynamic interface approach also allows each user group to view a level of detail suitable to its needs, a key requirement for data visualizations that will be used by multiple audiences.

The simplest and most straightforward approach, however, may be the one that directly solicits patient feedback in the design and refinement of data visualizations and overall information presentation. Though not used uniformly

across the field of patient education, the collection and analysis of user feedback has precedent within the online information environment. Within the realm of health information, the value of qualitative input from users has been understood and practiced for some time[35]. For example, patient feedback related to a Web site providing information about stereotactic ablative radiotherapy helped Web developers redesign the site such that patients were better able to understand scientific evidence, among other improvements[36]. Though some data visualizations may be more complex than the health information available on many consumer-focused Web sites, such studies indicate the importance of engaging users in design of data representations.

Conclusion

The development of forms of data representations that are more accessible to patients and their caregivers supports many goals providers and health care organizations are working to achieve. Visual data presentations can help patients enhance their understanding of their health, make more informed decisions about which they feel better in the short and long terms, contribute to improvements in quality of care and health outcomes (theirs and others'), and become active participants in learning health care systems. Visual data representations that are meaningful to patients, caregivers, and the health care professionals who care for them support patient engagement and facilitate patient-centered care, and their development should be a goal of the health care system.

References

1. Fox S. The social life of health information. Washington, DC: Pew Research Center. January 15, 2014. http://www.pewresearch.org/fact-tank/2014/01/15/the-social-life-of-health-information/
2. Fried TR. Shared decision making – finding the sweet spot. N Engl J Med. 2016; 374(2):104-106.
3. Institute of Medicine. Crossing the quality chasm: a new health system for the 21st century. Washington, DC: The National Academies Press, 2001. https://www.nap.edu/catalog/10027/crossing-the-quality-chasm-a-new-health-system-for-the
4. Wilbanks J, Friend SH. First, design for data sharing. Nat Biotechnol. 2016; 34(4):377-379.
5. Chia PL, Gedye C, Boutros PC, Wheatley-Price P, John T. Current and evolving methods to visualize biological data in cancer research. J Natl Cancer Inst. 2016; 108(8):djw031.
6. van der Heide I, Heijmans M, Schuit AJ, Uiters E, Rademakers J. Functional, interactive and critical health literacy: Varying relationships with control over care and number of GP visits. Pat Educ Couns. 2015; 98(8):998-1004.
7. Aldrich FK, Sheppard L. Graphicacy: the fourth "R"? Prim Sci Rev. 2000; 64:8-11.
8. Friel SN, Curcio FR, Bright GW. Making sense of graphs: critical factors influencing comprehension and instructional implications. J Res Math Educ. 2001; 32(2):124-158.
9. Kincaid JP, Braby R, Mears JE. Electronic authoring and delivery of technical information. J Instructional Development. 1988; 11:8-13.
10. Centers for Medicare and Medicaid Services. Using usability formulas: a cautionary note. Washington, DC: Centers for Medicare and Medicaid Services, 2016. Accessed September 2, 2016 at https://www.cms.gov/outreach-and-education/outreach/writtenmaterialstoolkit/downloads/toolkitpart07.pdf.
11. Alpert JM, Desens L, Krist AH, Aycock RA, Kreps GL. Measuring health literacy levels of a patient portal using the CDC's Clear Communications Index. Health Promot Pract. 2016. pii: 1524839916643703. [Epub ahead of print]
12. Brown SM, Culver JO, Osann KE, MacDonald DJ, Sand S, Thornton AA, Grant M, Bowen DJ, Metcalfe KA, Burke HB, Robson ME, Friedman S, Weitzel JN. Health literacy, numeracy, and interpretation of graphical breast cancer risk estimates. Patient Educ Couns. 2011; 83(1):92-98.
13. Nayak JG, Hartzler AL, Macleod LC, Izard JP, Dalkin BM, Gore JL. Relevance of graph literacy in the development of patient-centered communication tools. Patient Educ Couns. 2016; 99(3):448-454.
14. Lipkus IM. Numeric, verbal, and visual formats of conveying health risks: suggested best practices and future recommendations. Med Decis Making. 2007; 27:696-713.
15. Oudhoff JP, Timmermans DRM. The effect of different graphical and numerical likelihood formats on perception of likelihood and choice. Med Decis Making. 2015; 35(4):487-500.

16. Timmermans DRM, Ockhuysen-Vermey CF, Henneman L. Presenting health risk information in different formats: the effect on participants' cognitive and emotional evaluation and decisions. Pat Educ Couns. 2008; 73(3):443-447.
17. Ancker JS, Senathirajah Y, Kukafka R, Starren JB. Design features of graphs in health risk communication: a systematic review. J Am Med Inform Assoc. 2006; 13:608-618.
18. Henneman L, Oosterwijk JC, van Asperen CJ, Henko FH, Ockhuysen-Vermey CF, Kostense PJ, Claassen L, Timmermans DRM. The effectiveness of a graphical presentation in addition to a frequency format in the context of familial breast cancer risk communication: a multicenter controlled trial. BMC Med Inform Decis Mak. 2013; 13:55
19. Damman OC, De Jong A, Hibbard JH, Timmermans DRM. Making comparative performance information more comprehensible: an experimental evaluation of the impact of formats on consumer understanding. BMJ Qual Saf. 2015; pii: bmjqs-2015-004120 [Epub ahead of print]
20. Murphy M, Cavill N. Do physical activity infographics, report cards and policies influence behaviour? Sport Exerc Sci. 2015; 46:12.
21. Scott H, Fawkner S, Oliver C, Murray A. Why healthcare professionals should know a little about infographics. Br J Sports Med. 2016; 50(18):1104-1105.
22. Hawley ST, Zikmund-Fisher B, Ubel P, Jancovic A, Lucas T, Fagerlin A. The impact of the format of graphical presentation on health-related knowledge and treatment choices. Pat Educ Couns. 2008; 73:448-455.
23. Nilsen W, Kumar S, Shar A, Varoquiers C, Wiley T, Riley WT, Pavel M, Atienza AA. Advancing the science of mHealth. J Health Commun. 2012; 17(Suppl 1):5-10.
24. Bhat S, Ferraris A, Gupta D, Mozafarian M, DeBari VA, Gushway-Henry N, Gowda SP, Polos PG, Rubenstein M, Seidu H, Chokroverty S. Is there a clinical role for smartphone sleep apps? Comparison of sleep cycle detection by a smartphone application for polysomnography. J Clin Sleep Med. 2015; 11(7):709-715.
25. Chen J, Cade JE, Allman-Farinelli M. The most popular smartphone apps for weight loss: a quality assessment. JMIR Mhealth Uhealth. 2015; 3(4):e104.
26. Wilson H, Stoyanov SR, Gandabhai S, Baldwin A. The quality and accuracy of mobile apps to prevent driving after drinking alcohol. JMIR Mhealth Uhealth. 2016; 4(3):e98.
27. Duane M, Contreras A, Jensen ET, White A. The performance of fertility awareness-based method apps marketed to avoid pregnancy. J Am Board Fam Med. 2016; 29(4):508-511.
28. Schnall R, Rojas M, Bakken S, Brown III W, Carballo-Dieguez A, Carry M, Gelaude D, Mosley JP, Travers J. A user-centered model for designing consumer mobile health application (apps). J Biomed Inform. 2016; 60:243-251.
29. Gakumo CA, Raper JL, Cerice DK, Stand-Gravois MJ, Mugavero MJ. A qualitative study on health numeracy and patient-provider communication of laboratory numbers in older African Americans with HIV. J Assoc Nurses AIDS Care. 2016; pii: S1055-3290(16)30077-2 [Epub ahead of print]
30. Weissgerber TL, Milic NM, Winham SJ, Garovic VD. Beyond bar and line graphs: time for a new data presentation paradigm. PLoS Biol. 2015; 13(4):e1002128.
31. Saxon E. Beyond bar charts. BMC Biol. 2015; 13:60.
32. Weissgerber TL, Garovic VD, Savic M, Winham SJ, Milic NM. From static to interactive: transforming data visualization to improve transparency. PLoS Biol. 14(6):e1002484.
33. McCaffery KJ, Dixon A, Hayen A, Jensen J, Smith S, Simpson JM. The influence of graphic display format on the interpretations of quantitative risk information among adults with lower education and literacy: a randomized experimental study. Med Decis Making. 2012; 32:532-544.
34. Kopanitsa G, Veseli H, Yampolsky. Development, implementation and evaluation of an information model for archetype based user responsive medical data visualization. J Biomed Inform. 2015; 55:196-205.
35. Van den Haak M, van Hooijdonk C. Evaluating consumer health information Websites: the importance of collecting observational, user-driven data. In International Professional Communication Conference (IPCC). IEEE Xplore Digital Publishers; 2010:333-338.
36. Hopmans W, Damman OC, Timmermans DRM, Haasbeek CJA, Slotman BJ, Senan S. Communicating cancer treatment information using the Web: utilizing the patient's perspective in website development. BMC Med Inform Decis Mak. 2014; 14:116.

Clinically Relevant Filters to Consider When Designing A Visualization for Longitudinal Electronic Health Records

Filip J. Dabek, MSc, Jefferson E. McMillan, Jesus J. Caban, PhD [1]
National Intrepid Center of Excellence,
Walter Reed National Military Medical Center, Bethesda, MD

Abstract

Electronic Health Record (EHR) data has multiplied over the past decade, resulting in petabytes of longitudinal data available for analysis. Many visualizations and systems have been built to analyze this large, complex data, but few have been able to assist clinicians by supporting their analytical thought process and the methods in which they prefer to manipulate and filter a dataset. Through interactions and case studies with providers, we have identified and present in this paper the top filters that are clinically relevant and should be considered for when building visualizations that support longitudinal healthcare data.

1 Introduction

The adoption of EHR systems has multiplied over the past decade, leading to millions of patient records being converted to digital form and resulting in petabytes of data available for analysis. While this conversion has increased the availability of both patient and population level data, it has itself lead to an over saturation of information for healthcare providers [1–3]. In particular, providers now face the challenge of developing temporal understanding of their patients' histories and diagnoses.

While systems have been developed for analyzing large, complex clinical data, clinicians can become frustrated by those that do not allow them to customize their experience to fit their thought process. Compared to many other domains, clinical care providers have developed mental shortcuts and preferred strategies for analyzing their data [4]. Thus, an approach that is able to cater to clinicians' needs and simulate their thought process has the potential for wide spread adoption and improved analytical performance.

In this paper we present visualization filters developed through interactions and case studies with providers, which should be utilized when building visualizations and exploration systems. These filters allow clinicians to analyze different facets of a dataset, as well as reduce the amount of data visible for analysis so clinicians do not become burdened with the amount of information displayed. We provide an overview of existing systems and visualizations before presenting our filters in Section 3, showing how the effectiveness of visualizations and systems can be improved through the use of our clinically relevant filters.

2 Background

Many analytical systems built for large, complex data analysis were designed for use in domains other than healthcare and then later re-purposed to be used in the clinical setting. While these systems have shown promise in their original domains [2;5], they have not translated well to the clinical realm where providers require specially crafted analytical techniques and tools that cater to their unique clinical thought process. Specifically, basic event filters leveraged by these systems do not account for the large number of variables needed for clinical event analysis and do not account for the varying thought processes of different clinicians. Thus, these systems lack the flexibility necessary to be clinically applicable and lead to low adoption rate across the clinical spectrum.

Furthermore, there have been visualizations and systems that have been built specifically for the clinical domain, most notably CareFlow [6], DecisionFlow [7], OutFlow [8], and Frequence [9] in which the popular sankey diagram (example provided in Figure 1) has been utilized to visually display the various pathways and outcomes that a patient may undergo. While these systems have produced effective visualizations, they lack the power and flexibility for clinicians

[1] The views expressed in this paper are those of the authors and do not reflect the official policy of the Department of Army/Navy/Air Force, Department of Defense, or U.S. Government.

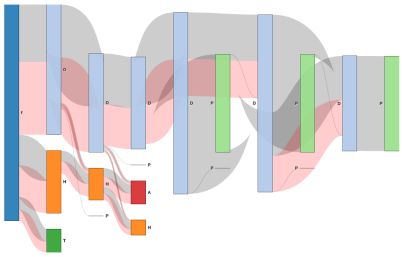

Figure 1: An example of a sankey diagram that is commonly used in visualization systems.

Figure 2: Chord diagrams are utilized to view relationships between a starting condition and potentially multiple end points.

Figure 3: A windrose is traditionally utilized in climatology, but can be leveraged in other domains to show relative frequency across ordered categories.

to filter and augment the dataset being displayed leading to an over-saturation of data and poor understanding of the underlying dataset.

In our own previous work, we too have developed a visualization system, VisXplore, for analyzing complex clinical datasets [10]. In VisXplore, the clinician is presented with various visualizations to be able to understand their data from a variety of angles. With this paper, we seek to expand our system and provide the clinician with even more control and customizability. In addition to the sankey diagram utilized by many existing visualization systems, our VisXplore system includes many different visualizations such as:

- **Chord:** as can be seen in Figure 2 where the clinician is able to analyze the correlation and relationship between various variables, events, and/or diagnoses.
- **Windrose:** as can be seen in Figure 3 where each patient can be plotted against eachother and the shapes of both plots can be utilized for a general overview of both patients.

3 Clinical Filters

While our system, VisXplore, and other various systems have embedded visualizations that were constructed to provide meaningful information to the user, clinicians require tools that allow them to manipulate and filter the dataset being displayed. Such functionality makes understanding and interpetation easier, thus allowing for greater analytical accuracy and adoption. A truly effective visualization systems needs to possess filtering tools that allow users to alter the dataset being displayed, as well as provide clinicians with the ability to reduce the amount of noise and only present information pertinent to that clinician. Therefore, through interactions and case studies with clinicians in which , we have identified a set of clinically relevant filters that allow a provider to analyze a dataset with the same mental shortcuts that they have developed over time. While the set of possible filters may be very large, we will present the top four filters that providers have communicated as being the most important and relevant to their analytical needs.

The primary task surrounding the filters that we will discuss revolves around analyzing a specific time interval of data, or what we refer to as "timeframe analysis". In Figure 4, we can see the filtering toolbox presented to clinicians with four tabs of interest: *Timeframe, Conditions, Dependencies*, and *Encounters*.

3.1 Timeframe

The timeframe tab was built to allow for providers to analyze only a specific "frame of time" for each patient, thus limiting the amount of data presented to only a subset of data available.

The content of this tab is shown in Figure 4 where it allows the provider to align their dataset on an event and select a specific time interval. Alignment on an event has been accomplished in other systems [5], but paired with the feature

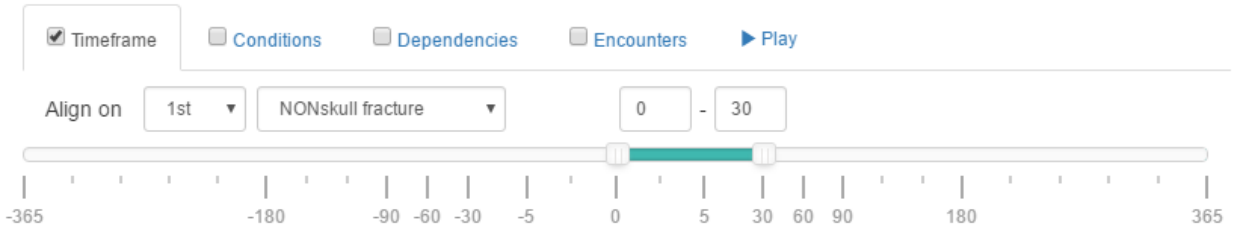

Figure 4: Our clinical-based filter toolbox. In the Timeframe tab, the clinician is able to align their dataset on a specific occurrence of an event in each sequence and then select the timeframe/time interval that they would like to analyze.

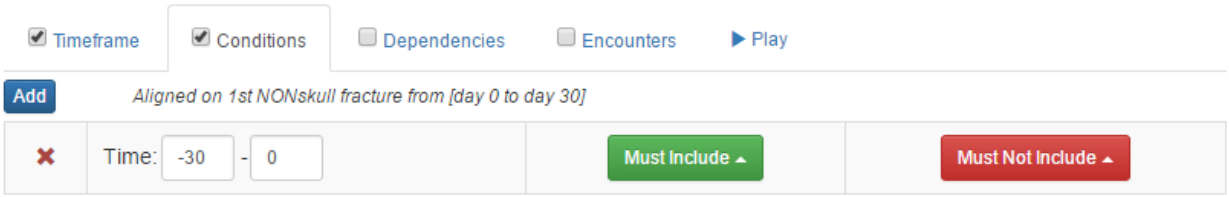

Figure 5: In the Condition tab, the clinician is able to add specific filters for multiple time intervals. In this specific example, the provider selected the events that must occur and must not occur in the 30 days prior to the first event of $NONskullfracture$ in the aligned dataset.

of selecting a time interval providers can analyze their data at different points in time. This method of filtering allows for users to look at various behaviors (i.e. only 30 days after the first time an event "A" occurred) and is particularly relevant if a provider is evaluating a patient and utilizes the visualization to better understand the patients potential path in a certain period of time.

3.2 Conditions

Next, Figure 5 contains the content in the Conditions tab where providers are able to define an unlimited number of conditions for the data. A condition is a constraint within a specific time interval defined by diagnoses that must be included or must not be included within the interval. When a clinician sets a condition, the dataset is filtered such that only sequences in the time interval that have or do not have certain diagnoses are shown. This method of filtering allows for only specific sections of the dataset to be filtered at a time as well as assist in cohort comparison. For example, a provider could be analyzing a dataset of patients that have developed cancer and then they could limit the display to only show those patients that suffered from headaches prior to being diagnosed with cancer. Furthermore, a provider could then perform cohort comparison to analyze the difference between the group of patients that suffered from headaches and those that did not. This filter provides multiple avenues for analysis and allows clinicians to customize based on their specific tasks (i.e. population analysis vs. single patient treatment).

3.3 Dependencies

The third filter in our toolbox is shown in Figure 6 where it allows the provider to apply an unlimited number of "dependency filters". Each dependency filter consists of the type of inclusion (can or must), the source event(s), and the target event; such that for each dependency the provider intends to only look at sequences where the patient possessed an encounter of either one (can) or all (must) source events prior to developing the target diagnosis. This filter is different from that of the conditions filter as it functions independent of time and provides a connection between two or more diagnoses and through the use of this filter, providers are able to identify patients that were diagnosed with certain diagnoses, regardless of the time filter, and then develop another diagnosis of interest.

As an example, for the case presented in Figure 6: $B, D \rightarrow PTSD$, the dependency filter would look at each

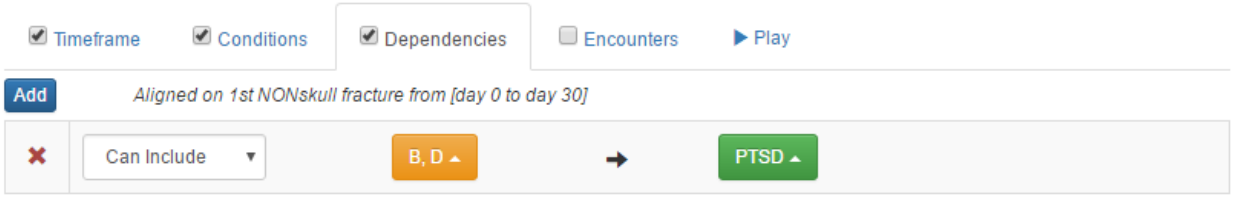

Figure 6: In the Dependencies tab, a clinician may filter based on events that supercede another event in a specific "dependency". In this example. provider selected that the diagnoses should contain $B, D \rightarrow PTSD$ where for each sequence in the dataset: either the event B or D must be in the sequence and $PTSD$ must exist after the first occurrences of B and D.

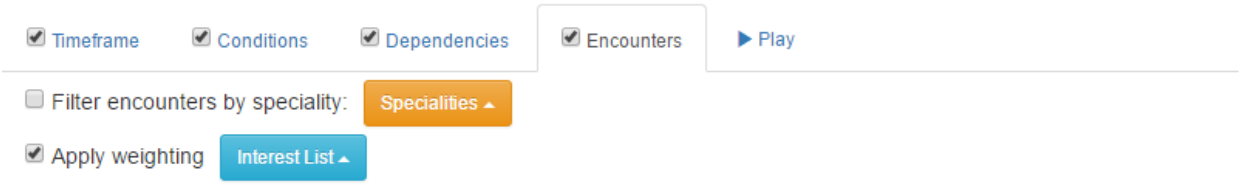

Figure 7: In the Encounters tab, the provider is able to only view encounters by the specialty that the encounter was classified under or identify the diagnoses that they are interested in to apply our clinical-based weighting formula.

sequence/patient in the dataset individually and identify the first occurrence of B and the first occurrence of D. Then it would identify which diagnosis occurred later and using its timestamp it would identify if the diagnosis $PTSD$ occurred after the timestamp. If the diagnosis $PTSD$ does not exist after the corresponding timestamp then the sequence would be removed from the dataset and not be presented for analysis. In addition, if both B and D do not exist in the sequence the sequence would also be removed. The filter has the ability to allow providers to specify different inclusion criteria including "can" or "must" include, where "can" corresponds to at least one source diagnosis must be present and "must" corresponds to all source diagnoses needing to be present. This additional functionality allows providers the flexibility to specify how rigid they desire the filter to be. Therefore, the dependency filter provides the flexibility needed to create a customized picture of a population (potentially similar to the one of a patient at bedside) and draw specific conclusions for this very specific cohort.

3.4 Encounters

The fourth filter incorporates a weighting formula described below which allows for dynamic adjustment to the graph, as shown in Figure 7. In this tab, the provider is able to both filter the type of encounters shown based on a particular specialty as well as turn on a specially designed weighting formula that alters the visualization being displayed to highlight the information that is pertinent to that provider.

Here, the provider is able to completely filter all encounters such that only the encounters performed by a provider of a specific specialty will be shown. This method allows for a provider to analyze only events of providers with specialties similar to their own and remove any diagnoses that may not be considered relevant. As an alternative, providers have the option of allowing the weighting algorithm, which will be described below, to highlight appropriate information, thereby causing the opacity of each part of the visualization to be set based on its relative importance. Both of these options serve to de-clutter the visualization and present providers with only relevant information needed for population understanding.

3.4.1 Weighting Formula

To draw attention to the important aspects of the graph, we sought to develop a weighting formula that could determine which aspects of a visualization should be highlighted for a specific provider. The formula was constructed by recog-

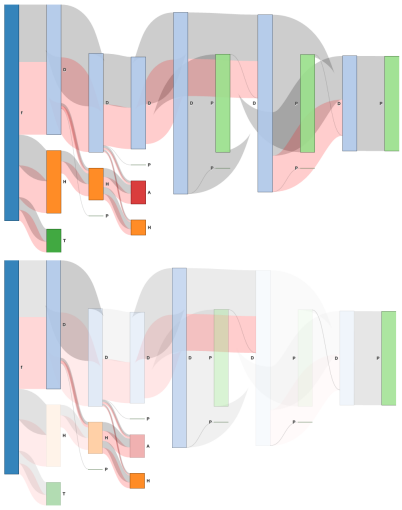

Figure 8: An example of a sankey diagram that is commonly used in visualization systems.

Figure 9: Chord diagrams are utilized to view relationships between a starting condition and potentially multiple end points.

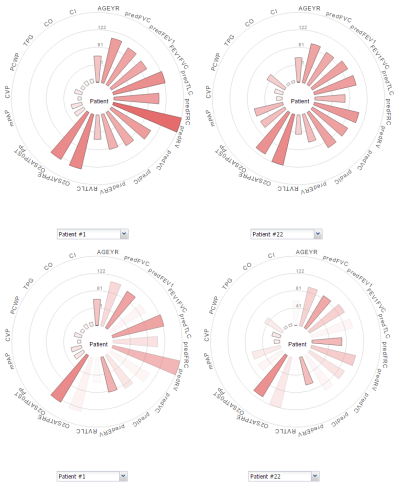

Figure 10: A windrose is traditionally utilized in climatology, but can be leveraged in other domains to show relative frequency across ordered categories.

nizing that, for longitudinal data, each provider is interested in: (a) diagnoses that are related to their discipline, (b) the temporal positioning of each diagnosis, and (c) the presence of certain conditions that a patient experienced. With this knowledge of provider interest, our formula is able to assist in highlighting the features of interest in a visualization by changing the opacity of each aspect of a visualization such that those that are less opaque (more transparent) are less important, while those that are more opaque (less transparent) are more important and of interest to the provider.

Prior to executing our formula, a provider would be required to provide a preference list in which they sort the diagnoses in order of their preference to analyze. Diagnoses at the beginning of the list are of greater interest and those at the end of the list are of less interest. However, when discussing this procedure with providers we recognized that requiring individual providers to do this at the initiation of the visualization would cause excessive burden and be an impediment to adoption. Thus, we surveyed providers from different specialties to understand how likely they were to utilize each diagnosis in their clinical decision making. Providers ranked each diagnosis/event type in order from most helpful to least helpful. The consensus rankings were stored for each specialty which could then be read from the provider's credentials each time the visualization was initialized.

While the provider does not need to specify a preference list as we built a consensus list for each specialty, the provider is required to indicate if there are any diagnoses that are of particular interest to them, such as if they are wanting to analyze all patients' pathways to developing PTSD. This could change between each time a provider accesses the visualization, thus the selection must be completed each time and can be seen in Figure 7 where the provider is presented with a dropdown for "Interest List".

With the preference and interest lists provided, we determined the opacity of each diagnosis through our weighting formula:

$$Opacity = 1 - \beta_1(R) - \beta_2(E) - \beta_3(PE) - \beta_4(P) \tag{1}$$

where R is the index of the diagnosis in the providers preference list (such that a lower index is more preferred), E is the binary indication of if the diagnosis is contained within the interest list, PE is the current diagnosis' distance from a node of interest (if it is not the node of interest), and P is the diagnosis' overall position in the sequence of events.

E, and thus PE, are derived from the provider input to the visualization; and P serves the purpose of adjusting all diagnoses to highlight the first diagnosis in the sequence.

While the weighting formula can be applied to any visualization, we show the effect of our weighting formula on the three visualizations that we showed in Section 2 in Figures 8, 9, and 10 where the original, unweighted visualizations are shown directly above the new visualization after applying our weighting formula. By comparing the weighted visualizations to the original visualiations, we can see that the provider's attention is now drawn to the nodes of interest in them. The amount of information is not overpowering, but all detail is available should a provider desire it. This allows for maximum flexibility of a visualization while maintaining its integrity and usability.

4 Conclusion

In this paper, we presented the top four clinically relevant filters that we have identified as being necessary for consideration when designing visualizations and visualization systems. These filters coincide with the thought process that clinicians follow as well as provide the ability to be able to manipulate and filter longitudinal EHR data effectively. With these filters, it can be seen that a visualization designed for use by clinicians requires direct input from the user to ensure that it follows their thought process. This shows that work from other domains and fields cannot be directly translated into the clinical setting, and that tools and techniques need to be specially constructed for clinical use.

With these filters that we have identified we look forward to integrating these filters into our VisXplore system and for the usefulness of these clinically relevant filters to be present across a variety of visualizations and systems in the clinical domain.

References

[1] Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1747–1756. ACM, 2011.

[2] Adam Perer and David Gotz. Data-driven exploration of care plans for patients. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 439–444. ACM, 2013.

[3] David Kindig and Greg Stoddart. What is population health? *American Journal of Public Health*, 93(3):380–383, 2003.

[4] Dan Maurino and Eduardo Salas, editors. *Human Factors in Aviation*. ElSevier, 2nd edition, 2009.

[5] Megan Monroe, Rongjian Lan, Hanseung Lee, Catherine Plaisant, and Ben Shneiderman. Temporal event sequence simplification. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2227–2236, 2013.

[6] Adam Perer and David Gotz. Visualizations to support patient-clinician communication of care.

[7] 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, December 2014.

[8] Krist Wongsuphasawat and David Gotz. Outflow: Visualizing patient flow by symptoms and outcome. In *IEEE VisWeek Workshop on Visual Analytics in Healthcare, Providence, Rhode Island, USA*, pages 25–28. American Medical Informatics Association, 2011.

[9] Adam Perer and Fei Wang. Frequence: interactive mining and visualization of temporal frequent event sequences. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*, pages 153–162. ACM, 2014.

[10] Filip Dabek and Jesus J Caban. Visxplore: flexible visualization system for analyzing complex clinical datasets. In *Proceedings of the 2013 Workshop on Visual Analytics in Healthcare*, page 14, 2013.

Evaluating Visual Analytics for Health Informatics Applications: A Progress Report from the AMIA VIS Working Group Task Force on Evaluation

David Gotz, PhD[1], David Borland[1], PhD, Jesus Caban, PhD[2], Dawn Dowding, PhD[3], Brian Fisher, PhD[4], Vadim Kagan, PhD[5], Danny T.Y. Wu, PhD[6]

[1]University of North Carolina, Chapel Hill, NC, USA; [2]Walter Reed National Military Medical Center, Bethesda, MD, USA; [3]Columbia University, New York, NY, USA; [4]Simon Fraser University, Surrey, BC, Canada; [5]SentiMetrix, Inc., Washington, DC, USA; [6]University of Cincinnati, OH, USA

Abstract

The American Medical Informatics Association (AMIA) Visual Analytics Working Group (VIS WG) established a Task Force on Evaluation (TFoE) in early 2016 to investigate the state-of-the-art in visual analytics evaluation and to provide a report documenting recommendations for visual analytics evaluation within the context of the medical informatics domain. This progress report documents the history of this task force, including its mandate and membership. This report also provides a brief summary of progress made so far, outlines future plans, and describes how additional members of the community can participate.

1. Introduction

The healthcare domain has long been a data-driven enterprise. From point-of-care decisions made by clinicians based on a patient's medical history, to longitudinal population studies that provide evidence for clinical practice guidelines (CPG), to individuals monitoring their own health through patient-generated health data (PGHD), the collection, organization, and utilization of information is at the center of nearly every aspect of modern medicine. This was the case in the era of paper charts, and continues as both the collection and utilization of data in medical practice has accelerated during the industry's shift toward a more digital and modern health IT infrastructure. For example, in the United States, the Office of the National Coordinator for Health Information Technology now reports that 96% of hospitals have an electronic health record system (EHR).[1] Reports suggest similar percentages of hospitals making progress toward meaningful use standards,[2] a set of criteria designed to access the capture and use of clinical data from EHR systems to improve quality, safety, and efficiency.

This ongoing digital transformation is producing large amounts of digital data, and is sparking a broad range of research and development aimed at enabling new data-driven methods for improving the healthcare system. One critical aspect of this wave of innovation has been in the design and development of effective ways to communicate data that can ultimately generate new knowledge and enable more insightful actions. Both the medical informatics and visualization research communities have recognized the growing importance of this challenge and have identified visual analytics as a critical area for technological innovation.[3,4] Visual analytics technologies support analytical reasoning about complex and large scale datasets using a combination of interactive visualization-based user interfaces and computational analysis. As such, these methods have the potential to help make data more interpretable and actionable for a range of healthcare user populations: from clinicians, to population health analysts, to patients, and to their caregivers and families.

However, despite the great promise of visual analytics to support more effective data analysis and decision making, it can be challenging to evaluate the benefits that a specific technology provides. This difficulty is recognized within the visualization community,[5] but is an even more critical hurdle in medical informatics applications where technologies must be rigorously proven before they can be widely adopted.

In early 2016, the American Medical Informatics Association (AMIA) Visual Analytics Working Group (VIS WG) established a Task Force on Evaluation (TFoE) to investigate the state-of-the-art in visual analytics evaluation and to provide a report documenting recommendations for visual analytics evaluation within the context of the medical informatics domain. This paper provides the history of the TFoE, describes its mandate and composition, and summarizes both its progress to date and future plans.

2. The Creation of a Task Force on Evaluation

Reflecting a growing interest in applying advances in visual analytics to the medical domain, the annual Visual Analytics in Healthcare Workshop[6] will be held for the seventh time this year. The workshop first took place in 2010 and has been held annually since then at either the AMIA Annual Symposium or the IEEE VIS Conference, reflecting the interdisciplinary nature of the topic. The emerging community fostered by this event was recognized by AMIA in 2015 with the establishment of the official VIS WG.

The VIS WG held its first annual meeting at the 2015 AMIA Annual Symposium. During the annual meeting, attendees were asked to suggest potential activities for the VIS WG to organize in its first year. One topic that resonated broadly during that discussion was the need to address best practices for evaluation of new technologies. The group recognized the both (1) the fundamental difficulty of evaluating visualization technologies, and (2) the critical importance of evaluation given the medical context of our work.

At the conclusion of those discussions, it was recommended that the VIS WG establish a task force to survey the state-of-the art in this area, and to recommend best practices for evaluation of visual analytics research within the medical informatics domain. The TFoE would be charged with developing a report to document its findings, with this article serving as an interim progress report.

Following the annual meeting, the VIS WG distributed a call for volunteers via both the AMIA VIS WG mailing list (restricted to AMIA members) and the VAHC email list[6] (representing a broader and more diverse community). All interested parties were invited to join the TFoE's first conference call on February 5th, and a total of 16 people called in to participate. Over subsequent month meetings, a group of seven people (all authors on this report) emerged as the core contributors to the task force: David Gotz (chair), David Borland, Jesus Caban, Dawn Dowding, Brian Fisher, Vadim Kagan, and Danny Wu. This team has broad representation, with members from industry, government, and academia.

3. Progress to Date

In this section we summarize the TFoE's preliminary results. Over the course of monthly meetings, beginning in February, the TFoE has engaged in two major threads of activity: an *interdisciplinary literature review*, and the *development of a framework* for characterizing evaluation methods specifically within the medical informatics domain.

3.1 Literature Review. We have identified three general domains that should be considered when studying evaluation techniques relevant to medical informatics: "traditional" visualization, health IT, and cognitive psychology.

Traditional Visualization. Evaluation in the visualization literature often involves user studies in which quantitative measures such as speed and accuracy are measured for specific visual representations. However, other evaluations techniques such as long-term case studies are also commonly employed. There are many examples in the literature discussing the unique challenges of visualization evaluation, the range of both quantitative and qualitative approaches that can be employed, and the relative strengths and weaknesses of those techniques.[5,7–10] Examples which apply some of these visualization evaluation methods within the medical informatics domain have also been described.[11,12]

Health IT. Within the health IT discipline, systems are typically viewed as comprised of several interacting components (e.g., the content of the system, the user interface, and the hardware on which an intervention is delivered), which are in turn implemented within larger equally complex systems (e.g., interacting health care organizations). This multi-layered systemic complexity makes the evaluation of health IT systems an enormously complex problem, with challenges including the identification of how the different components of the intervention (the health IT system) interact to produce outcomes, and the causal pathways or mechanisms by which they achieve those outcomes. The literature in this field has addressed these issues in various ways, many of which can be applied to visualization-based systems. For example, the Medical Research Council (MRC) framework for complex interventions provides an overview of the process by which an intervention can be developed and then evaluated.[13] More broadly, many have proposed models that consider ways to evaluate the effect of health IT system on outcomes, while taking into account the complexity of the context in which they are implemented.[14–21] Finally,

recent work exploring so-called “realist evaluation” methods have focused even more directly on evaluations based on understanding the interactions between (1) the context (the situation and factors where an intervention is implemented), (2) the mechanism (through which an intervention is thought to change behavior or other factors), and (3) the clinical outcome.[22,23]

Cognitive Psychology. One approach to dealing with the complexity of evaluation is to develop analytic methods that seek out regularities in cognitive task performance. These begin with aspects of human information processing that are said to be "architectural" in the sense that they are consistent across individuals and over time for a given individual. Many of these human capabilities, such as trichromacy, and scope of verbal short-term memory, are well-known to the human-computer interaction (HCI) community.[24,25] Others, such as the number and processing of attentional tokens (Pylyshyn's FINSTs[26]) are less commonly understood. As data displays become more complex and dynamic we may find that the psychological underpinnings of traditional HCI methods must be augmented by aspects of human cognitive architecture that are only now being investigated in the cognitive science and psychology communities.[27] One way to do this is for a cognitive psychologist to closely examine a video screen capture of the interface in use, looking for potential threats to human cognitive architecture. This draws from "close reading" methods used in the humanities, but is intended to understand the interaction of human cognitive architecture with the unusual perceptual situations generated by modern display environments. From this examination a laboratory study can be constructed that can evaluate whether those threats are real. For example, an examination of a proposed Next-Gen air traffic control interface generated a set of psychology studies that evaluated whether changes in viewer position in a moving-target display would adversely affect air traffic controllers' ability to track individual aircraft using the new interface approach.[28]

Not all cognitive task regularities are architectural. Many differ between individuals due to their individual capabilities. Laboratory studies of individual differences in performance may find patterns of behavior that are consistent for that individual but differ between individuals. Investigation of these patterns may lead to a "personal equation of interaction" which might enable an interface to be adapted to a given user's cognitive abilities as well as to their preferences.[29] As with the cognitive architecture work, this also generates quantitative measures, however in this case the tests are done entirely within subjects with an eye towards evaluating consistency of performance for a given individual in the task environment. Such methods may be especially germane to the medical informatics domain, in which many individuals with diverse backgrounds (e.g., patients, nurses, doctors) may interact with the same data in different ways.

Both cognitive architecture and individual difference studies fall along the X axis in Figure 1, *quantitative measurements*. While they may begin with examination of rich data (e.g. a screen capture video) the goal is to move to the laboratory for quantitative studies. If we are to address the Y axis of realism from a cognitive perspective we must find ways of building theory from rich data more directly. To address the qualitative Y axis in figure 1 we refer to the work of social scientists whose qualitative ethnographic research methods have been applied to examine organizational processes. Cognitive ethnography constitutes a special case in that it bridges ethnographic methods and cognitive task performance.[30] These approaches emerged from a new perspective in cognitive science that views cognition as a product of interaction of mental activities and information from the environment, often in the form of cognitive artifacts such as notation systems and visualization.[31,32] This labor-intensive method requires trained video analysts supported by software designed specifically for analysis of sociotechnical systems.[33]

The greatest challenge in our attempt to understand medical information systems lies along the diagonal in figure 1, where we examine how systems that include one or more human agents interact with the rich sensory environments that visual information systems can provide. While this is a new frontier, some progress is being made through the use of mixed-methods such as field experiments that manipulate some aspects of a complex task that is conducted in a realistic environment. Traditional social science methods such as grounded theory can be used here, and hybrid cognitive science approaches utilizing large-scale framework theories such as Clark's Joint Activity Theory are being developed.[34–38]

Effort on the literature review continues, however it has already proven fruitful in helping us form an organizational framework for evaluation.

3.2 Framework. The task force has developed and is continually improving a framework to organize the findings from the selected publications (Figure 1 (left)). In this two-dimensional framework, each publication can be

positioned based on the degree of its quantitative measures and its realism of settings and tasks. For example, a longitudinal study using highly quantitative measures will be placed on the top-right corner and regarded as an outcome study. The space can be divided into four regions, enabling the framework to categorize publications into four broad groups: *initial prototypes, task-based time and error studies, longitudinal case studies,* and *outcome studies*. These categories can be characterized based on (1) the degree of quantitative measurements and (2) the realism of the study tasks and environment. All four groups provide are valuable perspectives on evaluation and can give new insight about the frequency of studies that combine both quantitative rigor and realistic settings and tasks. The results from this study will be crucial for determining which areas of visual analytics in healthcare require more attention and are worthy of the investment in time and resources.

Figure 1(right) shows some of our preliminary results after reviewing 23 papers that introduce a visualization framework to explore clinical data. Each paper was reviewed and received a 0-5 score for the level of qualitative and quantitative evaluations that was performed. The size of the circles in Figure 1(right) represents the number of papers that received a specific score. Preliminary results show that a significant amount of papers describe a system and use some sort of qualitative measure to describe the benefits of the tool without providing detailed quantitative scores. It was encouraging to see that 21.7% of the papers have a balanced approach to describe and validate their frameworks as illustrated by the 2/2 and 4/4 scores.

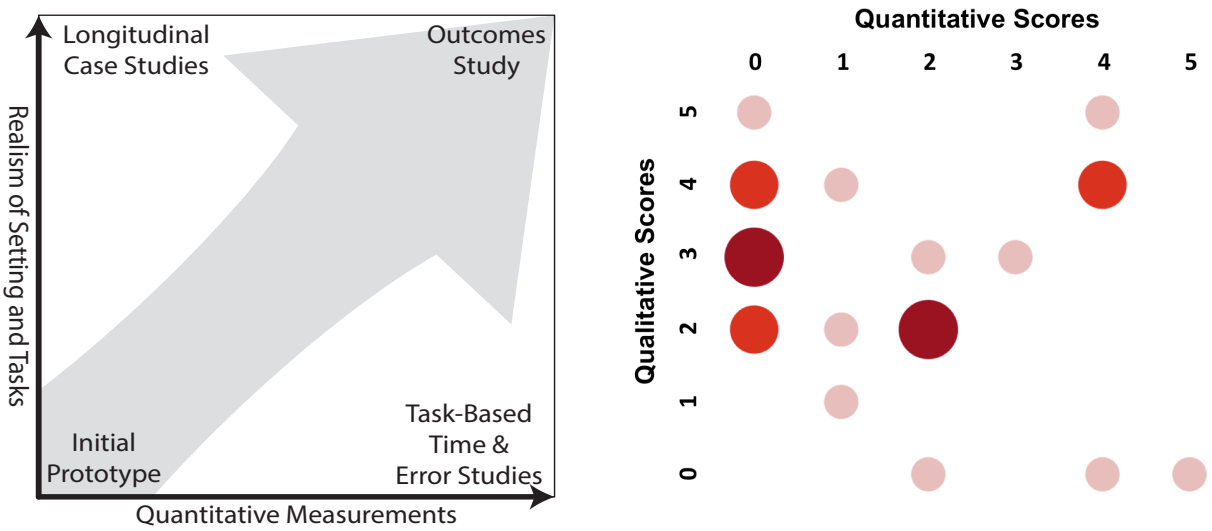

Figure 1. (left) A framework for characterizing the space of possible evaluations. The two key dimensions include the clinical realism of the evaluation (the Y axis) and the level of quantifiable evidence gathered (the X axis). (right) Preliminary results after reviewing 23 papers that introduce and validate a visualization framework to explore clinical data.

4. Future Agenda and Opportunities to Participate

In the first eight months of its existence, the AMIA VIS WG TFoE established the general framework for collecting and evaluating both existing and developing approaches for the visualization of health-related information. The efforts of the TFoE, however, are far from complete and there are many opportunities for members of the community to contribute to TFoE's efforts going forward.

One of the most critical items is the completion of the literature review–thus building the foundation for the remaining TFoE tasks. While the team has identified several publications in its early work, the universally accepted importance of visualization methodologies ensures that the universe of relevant publications is much larger. The core team would like to see participants from across different fields contribute to the growing collection and organization of relevant literature. Moreover, as we have started with our work on the interactive visualization of our

evaluation framework and related literature, possibilities exist for the development of tools to help gather, organize, and communicate TFoE findings. For example, there is the potential to explore automated methods which use modern data-mining platforms, such as Stanford Deep Dive,[40] to help identify new relevant publications as they become available as part of a dynamically updated repository.

Concurrently with the literature review, recommendations for best practices in terms of evaluation procedures must be developed in alignment with the framework being established by the Task Force. The creation of standard criteria and the corresponding guidelines for when certain methods are most appropriate will be an important step toward establishing a "gold standard" for evaluation activities when conducting visual analytics research in the healthcare domain.

In order to achieve wider visibility and to facilitate engagement of researchers and industry experts beyond the core community, the TFoE is planning to develop a public web site where up-to-date reports, tasks and challenges will be available for the general public to review. In addition to these documents, the website will contain the previously described interactive visualization of how existing identified literature fits within the proposed evaluation framework, and the planned best practices recommendations. This will be a critical tool in disseminating the results and collecting feedback from the community.

While this paper represents an update on the work in progress, a more comprehensive formal report covering the TFoE activities is planned for the future. However, as described above, much work remains to be done before such a report can be produced. All members of the broader VAHC community are invited to join the task force and to contribute to its ongoing work. Those interested in joining the TFoE, receiving notifications about future task force reports, providing feedback on reported results, or suggestions for future activities are encouraged to contact the TFoE chair David Gotz at gotz@unc.edu.

5. Conclusion

The AMIA VIS WG Task Force on Evaluation (TFoE) was established in early 2016 to investigate the state-of-the-art in visual analytics evaluation and to provide a report documenting recommendations for visual analytics evaluation within the context of the medical informatics domain. A team of seven experts have volunteered to work toward this goal, and this article serves as a progress report to the VIS WG community regarding the TFoE's progress to date. Progress includes an ongoing literature review and the development of a framework for characterizing different approaches to the evaluation process. The TFoE will continue these areas of work, with the goal of developing a final report for the VIS WG community in the coming months.

References

1. Adoption of Electronic Health Record Systems among U.S. Non-Federal Acute Care Hospitals: 2008-2015. Available at: http://dashboard.healthit.gov/evaluations/data-briefs/non-federal-acute-care-hospital-ehr-adoption-2008-2015.php. (Accessed: 1st August 2016)
2. Hospital Progress to Meaningful Use. Available at: http://dashboard.healthit.gov/quickstats/pages/FIG-Hospital-Progress-to-Meaningful-Use-by-size-practice-setting-area-type.php. (Accessed: 1st August 2016)
3. Caban, J. J. & Gotz, D. Visual analytics in healthcare – opportunities and research challenges. *J. Am. Med. Inform. Assoc.* **22,** 260–262 (2015).
4. Gotz, D. & Borland, D. Data-Driven Healthcare: Challenges and Opportunities for Interactive Visualization. *IEEE Comput. Graph. Appl.* **36,** 90–96 (2016).
5. Plaisant, C. The Challenge of Information Visualization Evaluation. in *Proceedings of the Working Conference on Advanced Visual Interfaces* 109–116 (ACM, 2004). doi:10.1145/989863.989880
6. Visual Analytics in Healthcare. Available at: http://www.visualanalyticshealthcare.org/. (Accessed: 1st August 2016)
7. Shneiderman, B. & Plaisant, C. Strategies for Evaluating Information Visualization Tools: Multi-dimensional In-depth Long-term Case Studies. in *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization* 1–7 (ACM, 2006). doi:10.1145/1168149.1168158
8. Elmqvist, N. & Yi, J. S. Patterns for visualization evaluation. *Inf. Vis.* 1473871613513228 (2013). doi:10.1177/1473871613513228
9. Plaisant, C., Grinstein, G. & Scholtz, J. Visual-Analytics Evaluation. *IEEE Comput. Graph. Appl.* **29,** 16–17 (2009).

10. Lam, H., Bertini, E., Isenberg, P., Plaisant, C. & Carpendale, S. Empirical Studies in Information Visualization: Seven Scenarios. *IEEE Trans. Vis. Comput. Graph.* **18,** 1520–1536 (2012).
11. Pieczkiewicz, D. S. & Finkelstein, S. M. Evaluating the decision accuracy and speed of clinical data visualizations. *J. Am. Med. Inform. Assoc. JAMIA* **17,** 178–181 (2010).
12. Sainfort, F., Jacko, J. A., McClellan, M. A. & Edwards, P. J. in *The Human-Computer Interaction Handbook* (ed. Jacko, Julie A.) 701–723 (CRC Press, 2012).
13. Craig, P. *et al.* Developing and evaluating complex interventions: the new Medical Research Council guidance. *BMJ* **337,** a1655 (2008).
14. Eisenstein, E. L., Lobach, D. F., Montgomery, P., Kawamoto, K. & Anstrom, K. J. Evaluating Implementation Fidelity in Health Information Technology Interventions. *AMIA. Annu. Symp. Proc.* **2007,** 211–215 (2007).
15. Takian, A., Petrakaki, D., Cornford, T., Sheikh, A. & Barber, N. Building a house on shifting sand: methodological considerations when evaluating the implementation and adoption of national electronic health record systems. *BMC Health Serv. Res.* **12,** 105 (2012).
16. Catwell, L. & Sheikh, A. Evaluating eHealth Interventions: The Need for Continuous Systemic Evaluation. *PLOS Med* **6,** e1000126 (2009).
17. Ammenwerth, E., Gräber, S., Herrmann, G., Bürkle, T. & König, J. Evaluation of health information systems-problems and challenges. *Int. J. Med. Inf.* **71,** 125–135 (2003).
18. Cresswell, K. M. & Sheikh, A. Undertaking sociotechnical evaluations of health information technologies. *Inform. Prim. Care* **21,** 78–83 (2014).
19. Ancker, J. S., Kern, L. M., Abramson, E. & Kaushal, R. The Triangle Model for evaluating the effect of health information technology on healthcare quality and safety. *J. Am. Med. Inform. Assoc. JAMIA* **19,** 61–65 (2012).
20. Eisenstein, E. L., Juzwishin, D., Kushniruk, A. W. & Nahm, M. Defining a framework for health information technology evaluation. *Stud. Health Technol. Inform.* **164,** 94–99 (2011).
21. Ammenwerth, E., Iller, C. & Mansmann, U. Can evaluation studies benefit from triangulation? A case study. *Int. J. Med. Inf.* **70,** 237–248 (2003).
22. Dalkin, S. M., Greenhalgh, J., Jones, D., Cunningham, B. & Lhussier, M. What's in a mechanism? Development of a key concept in realist evaluation. *Implement. Sci.* **10,** 49 (2015).
23. Randell, R., Greenhalgh, J. & Dowding, D. Using realist reviews to understand how health IT works, for whom, and in what circumstances. *J. Am. Med. Inform. Assoc.* **22,** e216–e217 (2015).
24. Pylyshyn, Z. W. The Role of Cognitive Architecture in Theories. *Archit. Intell.* 189 (1991).
25. Pirolli, P. Cognitive engineering models and cognitive architectures in human-computer interaction. *Handb. Appl. Cogn.* 441–477 (1999).
26. Pylyshyn, Z. W. *Things and Places: How the Mind Connects with the World*. (The MIT Press, 2007).
27. Fisher, B., Green, T. M. & Arias-Hernández, R. Visual Analytics as a Translational Cognitive Science. *Top. Cogn. Sci.* **3,** 609–625 (2011).
28. Liu, G. *et al.* Multiple-Object Tracking Is Based on Scene, Not Retinal, Coordinates. *J. Exp. Psychol. Hum. Percept. Perform.* **31,** 235–247 (2005).
29. Green, T. M., Jeong, D. H. & Fisher, B. Using Personality Factors to Predict Interface Learning Performance. in *2010 43rd Hawaii International Conference on System Sciences (HICSS)* 1–10 (2010). doi:10.1109/HICSS.2010.431
30. Hutchins, E. *Cognition in the Wild*. (A Bradford Book, 1996).
31. Hollan, J., Hutchins, E. & Kirsh, D. Distributed Cognition: Toward a New Foundation for Human-computer Interaction Research. *ACM Trans Comput-Hum Interact* **7,** 174–196 (2000).
32. Liu, Z., Nersessian, N. & Stasko, J. Distributed Cognition as a Theoretical Framework for Information Visualization. *IEEE Trans. Vis. Comput. Graph.* **14,** 1173–1180 (2008).
33. Fouse, A., Weibel, N., Hutchins, E. & Hollan, J. D. ChronoViz: A System for Supporting Navigation of Time-coded Data. in *CHI '11 Extended Abstracts on Human Factors in Computing Systems* 299–304 (ACM, 2011). doi:10.1145/1979742.1979706
34. Arias-Hernandez, R., Kaastra, L. T., Green, T. M. & Fisher, B. Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics. in *2011 44th Hawaii International Conference on System Sciences (HICSS)* 1–10 (2011). doi:10.1109/HICSS.2011.339
35. Clark, H. H. *Using Language*. (Cambridge University Press, 1996).
36. Bangerter, A. & Clark, H. H. Navigating joint projects with dialogue. *Cogn. Sci.* **27,** 195–225 (2003).
37. Clark, H. H. Coordinating with each other in a material world. *Discourse Stud.* **7,** 507–525 (2005).

38. Kaastra, L. T. & Fisher, B. Field Experiment Methodology for Pair Analytics. in *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization* 152–159 (ACM, 2014). doi:10.1145/2669557.2669572
39. Bostock, M., Ogievetsky, V. & Heer, J. D3 Data-Driven Documents. *IEEE Trans. Vis. Comput. Graph.* **17,** 2301–2309 (2011).
40. Shin, J. *et al.* Incremental Knowledge Base Construction Using DeepDive. *Proc VLDB Endow* **8,** 1310–1321 (2015).

POSTER PRESENTATIONS

Visualizing Edit Distances Between Kinase Inhibitor Names and English Words

Sandeep Jain[1], Jeremy L. Warner MD, MS[2-4]
[1]Vanderbilt University School of Medicine; [2]Division of Hematology/Oncology, Vanderbilt University; [3]Department of Biomedical Informatics, Vanderbilt University; [4]Vanderbilt-Ingram Cancer Center, Vanderbilt University Medical Center (all in Nashville, TN, USA)

Abstract: *Kinase inhibitors are promising chemotherapy drugs that are notoriously difficult to spell. Consequently, searching unstructured data such as Electronic Health Records (EHRs) for kinase inhibitor names is difficult and impedes research endeavors. We have created visual networks depicting the similarity of generic kinase inhibitor names amongst each other and across words in the English dictionary. This is an initial step towards a rational search query in EHRs and other unstructured text sources (e.g., social media platforms) for kinase inhibitor names.*

Introduction: *Kinase inhibitors* are an important class of chemotherapy medications that are used in many cancer treatment settings. A wealth of kinase inhibitor research data could be derived from Electronic Health Records (EHRs); however, searching for the names of kinase inhibitors within unstructured EHR data is difficult due to their problematic spellings. Kinase inhibitors' names are similar to one another, similar to various English words, and notoriously difficult to spell by clinicians as well as patients (e.g., erlotinib is misspelled elrontinib, etc.). These are all important considerations to keep in mind while searching for kinase inhibitor names in EHRs. To better and more intuitively understand the similarities between kinase inhibitor names to one another and to common English words, we have constructed a variety of network graphs to visualize word-to-word relationships. Complex multi-dimensional relationships such as these are often amenable to intuitive understanding when represented as graph visualizations. These visualizations provide an intuitive understanding of which kinase inhibitors will be most difficult to accurately capture from the patient EHRs due to similarities with other kinase inhibitors and/or with English words.

Methods: The generic names of 27 FDA approved small-molecule kinase inhibitors were used in this analysis. We used an electronically searchable version of the 2015 Merriam-Webster Dictionary as our source of English words for comparison. The computational linguistic concept of *edit distances* was used to measure the spelling similarities between words. Edit distance is the minimum number of operations required to turn one word into another. The specific type of edit distance we used is called the Levenshtein distance, which includes the removal, insertion, or substitution of a character in the word as viable operations [1]. The analysis and graphs were generated using the programming language R, version 3.2.4 (2016-03-10) [2]. The packages *stringdist* [3] and *igraph* [4] were used to compute edit distances and generate network visualizations, respectively. The package *wordcloud* was used to generate a novel hybrid layout [5]. Through *igraph* we were able to portray each vertex of the graph as a word, and each line (i.e., *edge*), as a representation of the edit distance. Two different force-directed network graph models were used in the creation of our visualizations. The models treat vertices on the graph as physical objects with attractive and/or repulsive forces on one another. The *Fruchterman-Reingold* model places attractive forces between connected vertices and repulsive forces between all vertices [6]. The *Kamada & Kawai* model creates spring forces between all pairs of vertices, utilizing Hooke's law [7].

Results: We have generated several network graphs showing the relationship between generic kinase inhibitor drug name spellings and dictionary word spellings (figures). Separate Fruchterman-Reingold and Kamada-Kawai layouts were created using identical word length and edit distance parameters (upper left and upper right). A novel "word cloud" hybrid layout was also created to uniquely visualize which generic drug names had the most dictionary words associated with them (lower left).

Discussion: The network graph visualizations created in this study show the large number of words that are similar to the generic drug names. Cognizant of these novel observations, we will better be able to modify EHR and free text searches of kinase inihibitor names. This will allow us to maximize the possibility of discovering drug exposures while avoiding false positives when using fuzzy search strategies. It is also possible that such word-to-word spelling analyses as demonstrated here could better guide ideal naming of drugs in the future.

VAHC 2016 Poster Submission

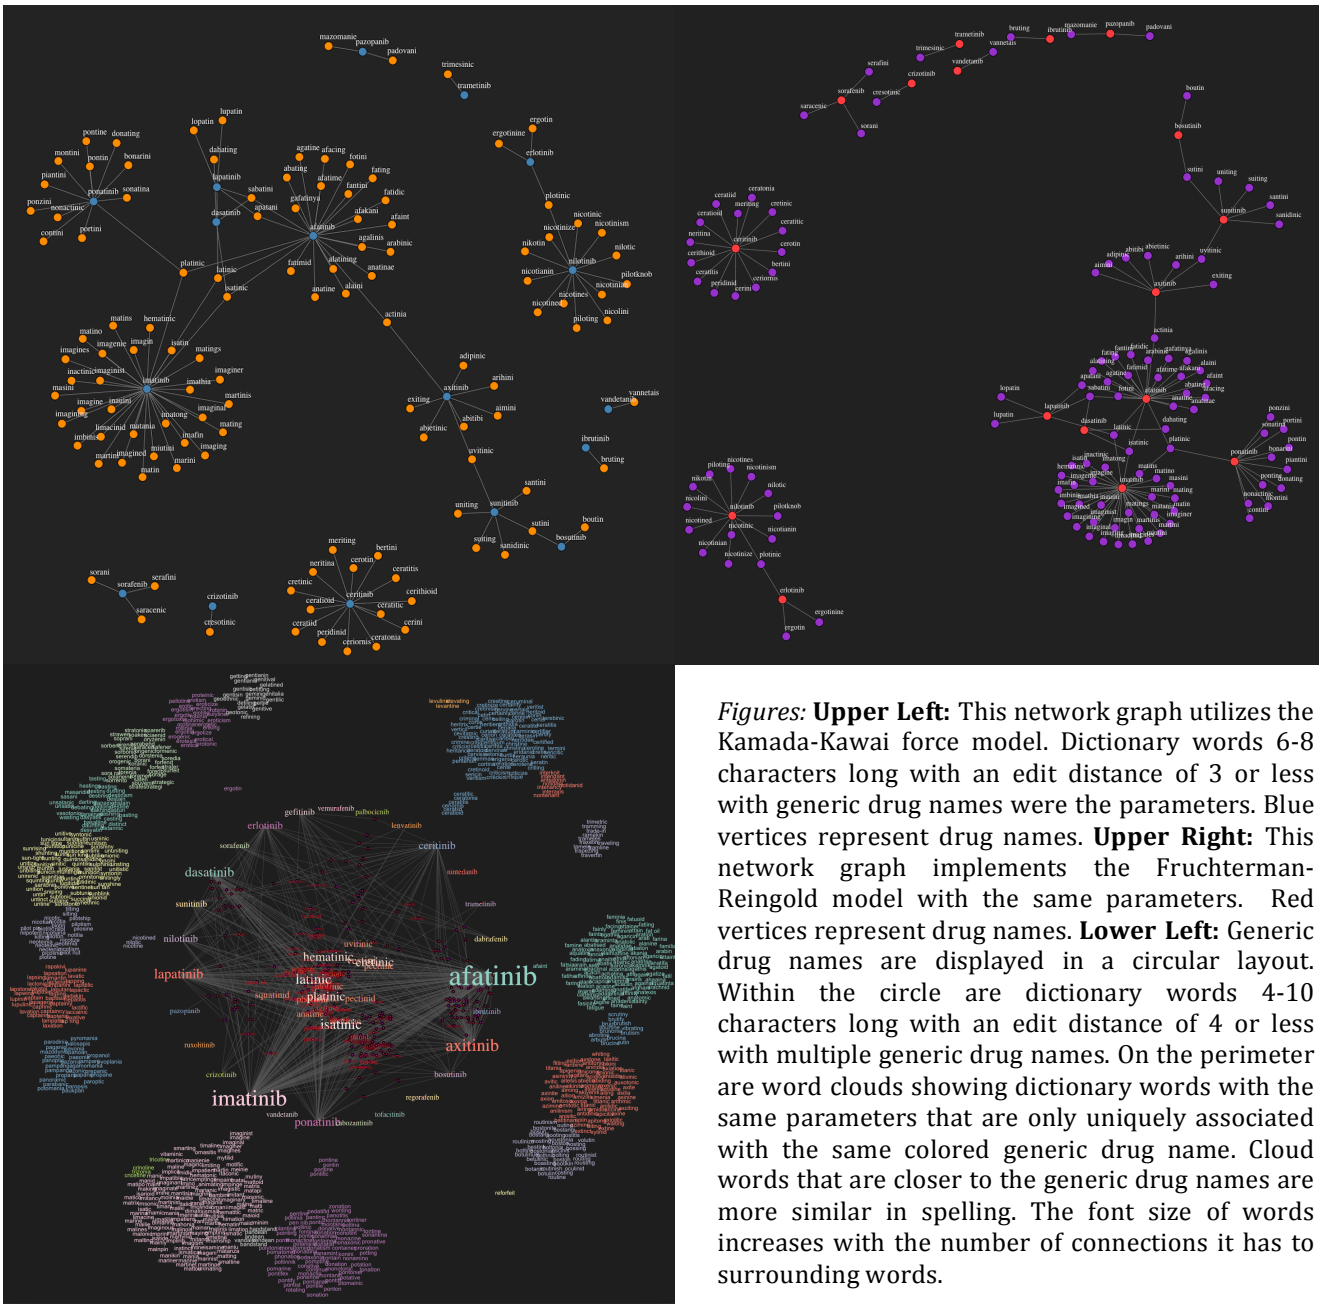

Figures: **Upper Left:** This network graph utilizes the Kamada-Kawai force model. Dictionary words 6-8 characters long with an edit distance of 3 or less with generic drug names were the parameters. Blue vertices represent drug names. **Upper Right:** This network graph implements the Fruchterman-Reingold model with the same parameters. Red vertices represent drug names. **Lower Left:** Generic drug names are displayed in a circular layout. Within the circle are dictionary words 4-10 characters long with an edit distance of 4 or less with multiple generic drug names. On the perimeter are word clouds showing dictionary words with the same parameters that are only uniquely associated with the same colored generic drug name. Cloud words that are closer to the generic drug names are more similar in spelling. The font size of words increases with the number of connections it has to surrounding words.

References:

[1] Levenshtein, Vladimir I. (February 1966). "Binary codes capable of correcting deletions, insertions, and reversals". *Soviet Physics Doklady*. **10** (8): 707–710.

[2] R Development Core Team (2008). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.

[3] van der Loo M (2014). "The stringdist package for approximate string matching." *The R Journal*, **6**, pp. 111-122. http://CRAN.R-project.org/package=stringdist.

[4] Csardi G, Nepusz T: The igraph software package for complex network research, InterJournal, Complex Systems 1695. 2006. http://igraph.org.

[5] Fellows, Ian. (2013). "Word Clouds." http://blog.fellstat.com/?cat=11.

[6] Fruchterman, T. M. J. and Reingold, E. M. (1991), Graph drawing by force-directed placement. Softw: Pract. Exper., 21: 1129–1164. doi:10.1002/spe.4380211102

[7] Kamada, T., Kawai, S. (1989): An Algorithm for Drawing General Undirected Graphs. Information Processing Letters, 31, 7-15.

Poster #2

CHAI: A Visual Interface for Examining Subject Matter Similarities across Intervention Chat Message Histories

Annie T. Chen[1], William R. Kearns[1], Emily F. Law[1,2], Nicole M. Alberts[1,2], Tonya M. Palermo[1,2]
[1]University of Washington School of Medicine; [2]Seattle Children's Research Institute

Abstract

Chat messages can involve a wide variety of topics, and the added dimension of time makes the problem of identifying similarities in chat message histories particularly complex. In this paper, we present Communication History Analysis Interface (CHAI), a visual interface to facilitate the identification of commonalities in chat message histories and the investigation of how these might be related to intervention outcomes.

Introduction

The analysis of time-oriented data has been identified as a particular challenge in visual analytics in health care, due to difficulties such as addressing the scale and complexity of data and intertwining patient conditions with treatment processes[1]. We are increasingly seeing examples of visual approaches to support analysis of high-dimensional temporal event data, including interactive multi-view visualizations and ad hoc statistical analytics[2], interactive visual cluster analysis[3], and ontology-based event sequence representation to support interactive exploration[4].

Intervention chat messages histories are one example of time-oriented data. They can be challenging to work with due to the complexity arising from high-dimensionality incurred through analyzing a variety of topics in sequence. As patients go through interventions, they may discuss many aspects of their experience that have bearing on their condition and/or the efficacy of the intervention. The dimension of time adds a layer of complexity to the problem of identifying similarities in these histories. Lastly, the narrative form of chat messages does not lend itself to the identification of topics in any given history, let alone the comparison of histories.

However, there are ways in which text mining, clustering, and visualization techniques may be combined to facilitate this process. In this paper, we present *Communication History Analysis Interface (CHAI),* a visual interface that is currently under development, which facilitates the identification of similar chat message histories, the discovery of common motifs, and the investigation of how these might relate to patient health outcomes.

Approach

We employed a three-fold approach: topic modeling, hierarchical clustering, and visual analysis.

1. Topic modeling

To determine whether there are similarities in the topics discussed, it was first necessary to identify the topics. To perform this task, we used a generative probabilistic modeling algorithm, Latent Dirichlet Allocation (LDA), which models documents as random mixtures over topics, where a topic is defined as a distribution of words[5]. Using the LDA implementation available within the MALLET toolkit[6], we identified the most common topic within each message. Then, we generated a sequence of topics by temporally aligning all messages for a given study ID.

2. Hierarchical clustering

We used the Smith-Waterman algorithm to determine the optimal alignment of the chat message topic sequences across study participants, and employed it as a similarity measure to assign the sequences to clusters using group-average agglomerative hierarchical clustering. The Smith-Waterman algorithm, often used to align and identify similarities between two biological sequences[7], is also employed here to facilitate the discovery of common motifs.

3. Visual Analysis

The interface offers multiple interactive views: Overview, Cluster Overview, and Cluster Detail. The Overview provides a high-level view of the data and a set of filters to explore subsets of interest. The other two views enable users to employ cluster analysis to identify study participants who may be similar to one another based on their message histories, in increasingly greater levels of detail. The Cluster Overview level displays the chat message histories and cluster summary statistics (e.g. cluster means on outcome variables) separated by cluster (Fig. 1).

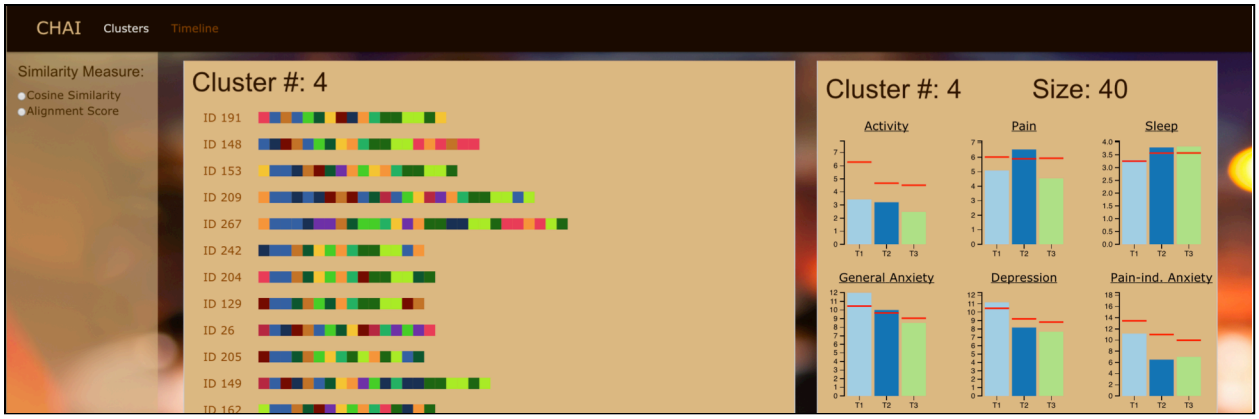

Figure 1. CHAI: Cluster Overview level.

In the Cluster Detail view, hovering over any given message reveals the corresponding message. The Cluster Detail view also features a legend that shows the high frequency keywords for each topic. A sample set of topics and keywords is depicted in Figure 2.

clues stress learning things move strategies working note pain drop find rest spend coach end people read beginning kinds

great sleep lifestyle work goal water habits making sounds move good set goals make day specific activity change working

skills pain learned ability great job looked anytime move made efforts proud improved encourage improve activity good check work

Figure 2. Sample topic keywords.

Case Study of an Analysis of Chat Message Histories for an Online Chronic Pain Intervention

We illustrate the potential use of CHAI to analyze chat message data from an Internet-delivered cognitive-behavioral therapy (CBT) intervention (Web-based Management of Adolescent Pain; Web-MAP) for youth with chronic pain and their parents[8]. We recently demonstrated the benefits of Web-MAP compared to an online education control condition in reducing activity limitations, improving sleep quality, improving parenting behaviors and parent-perceived impact. The dataset is comprised of chat message histories of participants randomized to the Web-MAP arm of the trial and includes messages between two types of dyads: online coach-youth, and online coach-parent. Online coaches initiated asynchronous email messages to participants in response to completed homework assignments through a secure message center. Online coach messages were guided by a structured online coaching manual, and focused on praising homework completion and problem solving barriers to skills practice. Participants had the option to respond to these messages or initiate new messages through the message center.

In order to identify similarities in chat message histories, we might access the Cluster Overview and Cluster Detail views. If we examine Cluster 4 in the Cluster Overview (Fig. 1), we see that, at the beginnings of the message histories, the blue topic is extremely common. Examining these messages reveals that they are the coaches talking with parents about noticing clues that indicate that their children's stress levels are increasing. Towards the end the message histories, we see discussions concerning lifestyle changes (dark green) and coaches reinforcing parents' work with their children (light green). Examining the common motifs of each cluster, along with the summary statistics of treatment outcomes of the cluster, can lead to insights concerning how aspects of the treatment experience are associated with treatment outcomes for the corresponding participants.

Conclusion

This paper described CHAI, an interface that is being developed to support exploration of similarities in chat message histories. The interface facilitates visual identification of similarities using an approach that combines text mining, cluster analysis, and multiple interactive views. A particular strength of this application is that it leverages text mining and cluster analysis techniques to identify automatically identify thematic similarities in narrative data that would otherwise be difficult to detect through manual examination, but leaves the difficult work of interpretation of patterns of significance in the hands of domain experts through interactive visual analytics features. Current development efforts include the implementation of additional clustering methods and on-demand statistical analysis to extend the analytic support provided by the application.

References

1. Aigner W, Federico P, Gschwandtner T, Miksch S, Rind A, Caban JJ, et al. Challenges of Time-oriented Data in Visual Analytics for Healthcare. Proc IEEE [Internet]. 2012 [cited 2014 Oct 5]; Available from: http://publik.tuwien.ac.at/files/PubDat_210197.pdf

2. Gotz D, Stavropoulos H. Decisionflow: Visual analytics for high-dimensional temporal event sequence data. IEEE Trans Vis Comput Graph. 2014;20(12):1783–92.

3. Gotz D, Sun J, Cao N, Ebadollahi S. Visual Cluster Analysis in Support of Clinical Decision Intelligence. AMIA Annu Symp Proc. 2011;2011:481–90.

4. Tao C, Wongsuphasawat K, Clark K, Plaisant C, Shneiderman B, Chute CG. Towards event sequence representation, reasoning and visualization for EHR data. In: Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium [Internet]. ACM; 2012 [cited 2015 Nov 1]. p. 801–6. Available from: http://dl.acm.org/citation.cfm?id=2110461

5. Blei DM, Ng AY, Jordan MI. Latent dirichlet allocation. J Mach Learn Res. 2003;3:993–1022.

6. McCallum AK. MALLET: A Machine Learning for Language Toolkit. 2002.

7. Smith TF, Waterman MS. Identification of common molecular subsequences. J Mol Biol. 1981 Mar 25;147(1):195–7.

8. Palermo TM, Law EF, Fales J, Bromberg MH, Jessen-Fiddick T, Tai G. Internet-delivered cognitive-behavioral treatment for adolescents with chronic pain and their parents: A randomized controlled multicenter trial. *Pain* 2016;*157*:174-185.

2016 Workshop on Visual Analytics in Healthcare
Poster #3

An Interactive Tool for Clinical Data Mining and Visualization

Qianxi Li[1], Yugang Jia[1], Merlijn Sevenster[1], Laura J. Olivieri[2], Craig A. Sable[2], Yue-Hin Loke[2], Justin Wiggs[2],
[1]Philips Research North America, Cambridge, MA, USA, [2]Children's National Health System, Washington, DC, USA

Abstract - In the era of value based healthcare, cardiologists rely on clinical data mining or visualization tools to perform data-driven studies such as effectiveness or efficiency measurements and cohort analysis [1]. Challenges for providing such tools include making available of necessary data for analytics, complexity and lengthiness of cohort selection, visualization of data and results, etc.. In this paper, we present an interactive tool that integrates multiple data sources (ECG, ECHO and MRI) to assist clinical researchers with data mining and visualization. One key component is to provide an interactive cohort selection dashboard to assist cohort selection. It provides a graphical way to present the distribution of cohorts and also provides real-time feedback when user changes cohort definitions. Tested with two use cases, it has been demonstrated that our tool can expedite the process of data mining. Using Tableau [2], visualization of data and analysis results is enabled in a manner that is meaningful to clinical researchers. On-demand analysis, in ad hoc fashion is also supported.

I. Introduction

In the era of value based healthcare, there is a need to provide clinical researchers with data mining or visualization tools to perform data-driven studies for a variety of topics, including cohort analysis, effectiveness or performance measurements, across multiple clinical databases [1]. Challenges for providing such tools include making available of necessary data for analytics, cohort selection, visualization of data and results, etc.. Clinical researchers are oftentimes practicing physicians who did not receive education in relevant technical areas, such tooling also needs to be comprehensible to the lay user.

While huge volume of clinical data is routinely collected as part of the care process, mining thereof remains challenging especially if the data is stored in disparate databases. Since the necessary data for analytics is stored in disparate databases, integration across multiple databases is therefore required. Also, clinical databases are transactional databases, which were never designed for data querying. Thus research queries on such databases are very complex, which creates a barrier to data access. This also raises a need to set up a data warehouse infrastructure that natively encompasses analytics data model to support efficient and transparent data querying and interactions. Furthermore, inconsistency of the same data fields across multiple data sources or a variety of data formats requires pre-processing of data.

Cohort selection is another important yet difficult step to support data analytics. Cohort selection is typically a complex and laborious process, because cohort definition in a formal query language is complicated, especially if the definitions evolve in the course of the analysis after inspection of cohort distribution. Without a clear understanding of data distribution in the cohort, it is typical that clinical researchers will go through a few iterations of cohort definition revising process before identifying a cohort to support further analysis.

Providing visualization tools also plays an important role in facilitating data mining. Although tools, such as Tableau [2], or D3.js exist, data mining will not be enabled if we fail to pre-process the data and set up a suitable data warehouse infrastructure for reasons discussed above.

In this paper, we present an interactive tool allowing clinical researchers to perform data mining across multiple cardiology databases (ECG, ECHO and MRI). An interactive visualized dashboard has been incorporated to assist cohort selection. Assisted by Tableau, it allows visualization of any data or analysis result. The dashboard has been tested with two use cases and the results demonstrate that it is scalable to support a wider range of clinical analytics topics and it helps to expedite the process of data mining.

II. Methods

The system consists of five integrated components, see in Figure 1.

Data Pre-processing Engine: This engine extracts and transforms data from multiple data sources and loads the data into a data warehouse.

Data Warehouse: The infrastructure that stores data using a comprehensive analytic data model to support efficient and transparent querying of data for analysis. With each added data source, a flat table is created to store the pre-processed data.

Cohort Selection Engine: This engine prepares the initial cohort selection table by including all the population. Then, by applying filters and temporal restriction, clinical researchers can select the cohort of interest.

Cohort Distribution Dashboard: This dashboard helps clinical researchers to interactively investigate the data distribution of cohorts.

Interactive Result/Data Visualization: clinical researchers can visualize the result or data, using Tableau.

III. Results

We have integrated multiple cardiology databases from a US-based academic hospital, containing ECG, ECHO and MRI interpretation data.

The system supported a wide range of analytics topics, including quality measurement, effectiveness evaluation of

procedure or test, cohort analysis and performance measurement. We use one of the use cases to demonstrate our result. In this particular case, the clinical researcher wants to test the hypothesis that the same measurement collected from MRI scan and ECHO test for the same patient could be different, especially in the abnormal cohort.

Data extracted and transformed from source MRI and ECHO databases will be stored in the data warehouse (DW). A cohort table in DW is generated by applying filters, in Tableau or DW, according to a definition for abnormal versus normal patients, see Table 1.

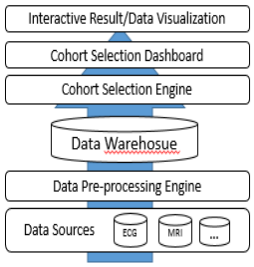

Fig 1. System Architect

Abnormal	MRI: Diagnostic statement indicates aortic valve insufficiency or aortic valve regurgitant > 20% or mitral valve regurgitant > 20% ECHO: Diagnosed with arotic valve insufficienty or mitrial valve insufficiency Find the study pairs in which MRI and ECHO tests are peformed within 6 months
Normal	MRI: Diagnosed with LV size normal and LV EVDI Zscore within (-2,2) ECHO: Diagnosed with LV size normal Find the study pairs in which MRI and ECHO tests are peformed within 6 months

Table 1 Cohort Definition

We have designed an interactive cohort selection dashboard, see Figure 3 that facilitates clinical researchers to understand the distribution of data within each cohort. Graph 'MRI' indicates that in the database, we have identified 82 abnormal MRI studies and 519 normal MRI studies. Within 3 months before MRI studies, 297 (56 + 241) cross referenced ECHO studies are identified, among which 241 cases are normal. The vertical axis of the three graphs in the upper area indicates the number of records found. The horizontal axis corresponds to the timeline. The MRI study sits at position 0 because we are measuring the time difference of ECHO studies compared to MRI studies, using MRI studies to cross reference studies in ECHO database. Studies in the class '3 months before MRI' indicates that the time difference between the MRI and cross-referenced ECHO studies is less than 3 month and the ECHO test is performed before the MRI scan. The temporal filter allows physicians to alter the timeline. For example, if only '-1, -2, -3' are selected, in the graph 'ECHO – Before MRI', only one column, '3 months before MRI' will be displayed. Other related fields will also be updated accordingly. The statistical summary table indicates total number of samples identified. By applying the filter 'ECHO_Cohort' or 'MRI_ABNORMAL', physicians can select to see the cohort of interest. The pie charts indicate the number of records identified in each cohort using specific criteria. For instance, in 'Details of Distribution - MRI', out of total of 82 abnormal MRI, 50 are identified using criteria 'Diagnostic statement indicates aortic valve insufficiency'. Enabling interactive investigation of cohort distribution provides a clearer view of the cohort, assisting clinical researchers in revising or finalizing cohort definition.

In the case when clinical researchers want to revise the definition of cohort, the definition could be updated using SQL like language or by applying filters in Tableau. For example, by unchecking 'MV regurgitant abnormal' in filter 'MRI Abnormal Definition' in Figure 3, definition for MRI abnormal cohort will be updated to exclude this particular criteria. Upon updating, the dashboard will automatically reflect the changes to provide real-time feedback.

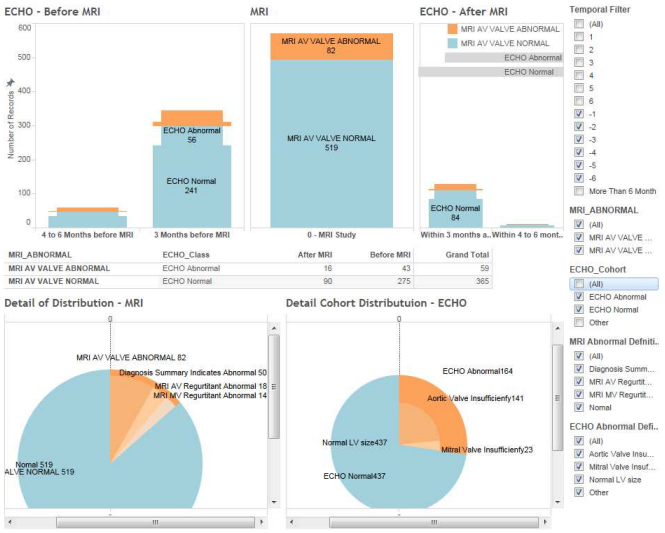

Fig 3. Cohort Distribution Dashboard

Ad hoc analysis is also enabled, since the necessary data for analytics has been made available and Tableau offers self-service environment for clinical researches to access and visualize data, see Figure 4.

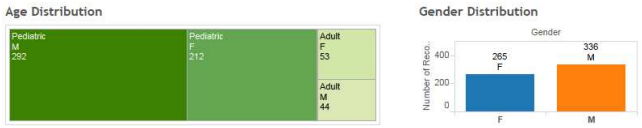

Fig 4. Ad hoc Analysis to Display Age and Gender Distribution within Cohort

It is displayed in Figure 5 that the system enables visualization of results, in various ways and in an interactive fashion.

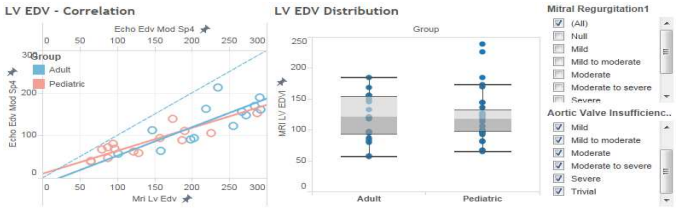

Fig 5. Visualization of analytics results

IV. Conclusion

The system has been successfully used at a US-based academic medical center for testing clinically impactful hypothesis, which results in clinical research abstracts. Results shows that the tool has enabled data mining across multiple cardiology data sources (ECG, ECHO, MRI). It facilitates cohort selection. It allows visualization of results in an interactive way. And it has greatly expedite the process of data mining.

References

[1] Jacqueline Weinberg, Puran Bhagat, Szymon Bieganski, A. Shane Brown, Melvin Kanasseril, Tim Lavers, Kirk Spencer, Craig Sable. Comparison of Left Ventricular Hypertrophy by Electrocardiogram and Echocardiography Using Novel Data Analytics Tool. National Conference & Exhibition, 2015

[2] http://www.tableau.com/

Poster #4
2016 Workshop on Visual Analytics in Healthcare

Genetics-based Motivation Viz ("GMV"): Visualizing Direct-to-Consumer Genetic Test Results To Empower Health Behavior Change

Kenyon Crowley[1,2,3,4], Joohee Choi[1], Rohan Bondili[1]
[1]University of Maryland iSchool, College Park, MD; [2]Human-Computer Interaction Lab (HCIL), College Park, MD;
[3]Robert H. Smith School of Business, College Park, MD;
[4]Center for Health Information & Decision Systems (CHIDS), College Park, MD

ABSTRACT
The human genome is a complex mix of approximately 10 million SNPs (single nucleotide polymorphisms). Each SNP represents a difference in a single DNA building block and may be correlated with traits and health conditions. The translation of complex multidimensional genetic data into actionable health behavior information to prevent disease is a difficult task. Our team developed the Genetics-based Motivation Viz ("GMV") in order to help individuals and health practitioners identify, motivate and prioritize genetics-based health behavior change opportunities. GMV incorporates an abstraction of a validated health behavior change framework, in combination with an individual's direct-to-consumer genetic testing (DTC-GT) results (including their health condition risk, moderating risk factors, strength of evidence supporting claim), juxtaposed with their self-reported ease of doing certain health behaviors (such as eating a low fat diet). GMV was deployed in a web browser and evaluated via a survey with 16 respondents. The majority of respondents indicated GMV was at least somewhat useful and understandable, but also identified areas for future work. With further development, GMV may be effective in translating genomic data into meaningful health behavior change information useful for individuals and health practitioners.

INTRODUCTION
The use of genetic information by consumers and health practitioners is growing rapidly and is expected to take on added significance in the diagnosis, prevention and treatment of disease in coming years.[1] Direct-to-consumer genetic testing (DTC-GT) refers to testing sold directly to consumers via the Internet, television, or other marketing venues without involving health care professionals. In general, DTC-GT utilizes mitochondrial DNA (maternal contribution), Y-chromosome (paternal contribution), and markers on autosomes (ancestral information) to provide data and interpretation about one's genetic makeup and potential consequences of this makeup (Ibid). Consumers' motivations for genetic testing typically include: general curiosity; improving their general health; ascertaining the risk of a particular condition; or, planning for future children.[2,3] Disease risk is a complex web of genetic, behavioral and environmental factors. Individual differences, including sociodemographic and psychological attributes as well as one's personal experiences (e.g. family history and first-hand experience with a disease) may influence one's genetic risk perceptions, and the use of genetic information.[4]

DTC-GT supporters believe that the information these tests provide may lead to better understanding of one's health, more empowered decision-making, and motivate adoption of beneficial health behaviors. [2,3] DTC-GT skeptics cite issues with the validity of testing, interpretability of results and potential unwarranted health services use from under-informed consumers.[1,5] When individuals are told they do not carry genetic-risk variants, the possibility they will interpret results to mean that they do not need to reduce behavioral risk factors is also a concern.[6] The United States Food & Drug Administration has barred certain DTC-GT service providers from providing health results, citing insufficient evidence to support genetic test report health claims.[7]

The current evidence is mixed regarding a consumer's health behavior change following a genetic screening.[8] McGuire finds that a sizable percentage (78%) of patients need help interpreting personal genomic testing results, which the GMV aims to do.[9] The primary aim of the GMV is to support the following users/needs: (a) A wellness coach or healthcare provider ("practitioners") seeking to understand a client's genetic disease predispositions and advise his/her client on behaviors to take to reduce the risk of genetic-based disease; and (b) An individual seeking to make decisions about behaviors to potentially reduce their risk of disease.

RELATED WORK
Our research is informed by the literature regarding: health behavior change; genetic information use by consumers, especially direct-to-consumer use; and, visualizing genomic data.

Chronic diseases including heart disease, cancer, lung diseases and diabetes are the most frequent causes of death in the United States.[10] Behavioral factors, particularly tobacco use, diet and activity patterns, alcohol consumption and avoidable injuries are among the most prominent contributors to mortality.[11,12] However, with changes to certain behaviors, individuals can dramatically reduce the severity and likelihood of disease. However, changing behaviors is difficult. Models to foster health behavior change have applied numerous techniques with varying success.[13] Fogg has a highly regarded and validated behavior change model, which finds successful behavioral change action is a function of the individual's motivation, their ability to perform a behavior and external triggers.[14] We tailor this model towards our DTC-GT use case and GMV.

An individual's perceived seriousness and controllability of the genetic test- identified disease is a major health behavior-moderating factor.[15] Other general factors that affect consumer's perceptions of genetic and pharmacogenetic testing include the clinical, epidemiological and economic evidence related to the information they must process.[16] These moderating factors are leveraged in GMV.

Determining how best to integrate and visualize genetics-related disparate data types to better understand biological systems is hard. Genetic viz tools need to combine diverse data forms, such as clinical information together with genomic data. These tools require visual representations that scale efficiently to thousands or millions of elements.[17] There are many stand-alone and web-based genome viz tools, which focus on helping scientists to explore, interpret and manipulate their data.[17,18,19] However, consumer tools are limited, and no genetic viewer for a health coach has been found in our literature search. Recent work advocates for more customizable interfaces for consumer genetic-data information processing using commercially available tools, such as Tableau.[20]

Figure 1. GMV (plus notations)

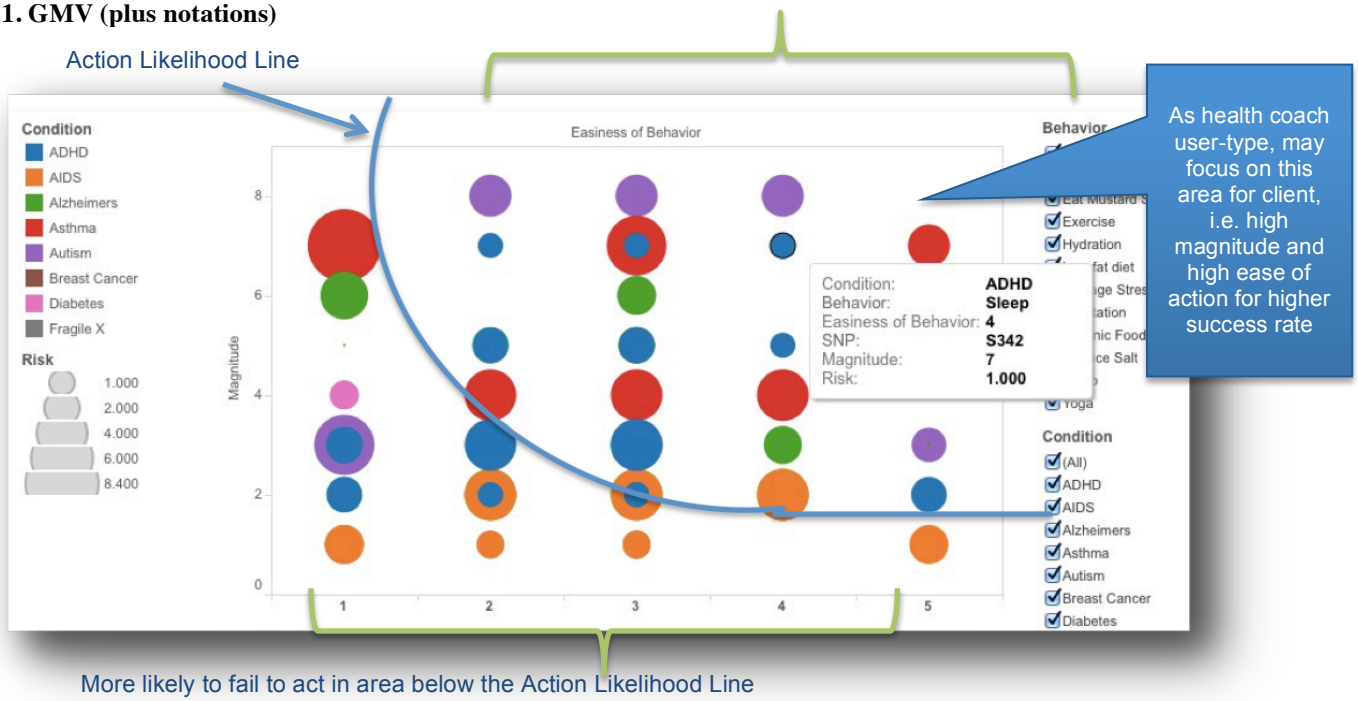

RESULTS

GMV was created using Tableau v9.1 software. DTC-GT raw data obtained via 23andMe was processed through the Promethease service (raw DNA data from multiple labs may be used), which provides SNP-specific disease risk, magnitude of importance, relevant risk-moderating behaviors based on the scientific literature in SNPedia. GMV combines this data file with a consumer's self-reported ease in performing certain health behaviors (such as eating a low fat diet) that may lessen their risk of contracting certain health conditions (such as heart disease). Figure 1 shows the GMV as it is displayed on Tableau Public at http://go.umd.edu/GMV.

Our GMV idiom uses color, size, and position to communicate genetic information. The GMV Y-axis shows the magnitude of an individual's SNPs. Magnitude is a subjective measure of interest generated by the SNPedia community, and ranges from 0: common genotype for which nothing interesting is known, 3: probably worth your time, to 10: really significant information (such as strong evidence of SNP's health impact). The GMV X-axis shows ease of behavior as reported by the individual, ranging from 1:not easy to 5: very easy). Node color matches the health condition associated with the SNP. Health risk of SNP is represented by node size, with a larger node meaning greater risk. Rollover the node and a pop-up provides text details about the disease condition, behavior, ease of behavior, SNP name, magnitude and risk variables. Filters are provided that can reduce the information density by disease condition and behaviors.

An interpretation example: A BRCA1-associated SNP (e.g. rs28897696; BRCA is associated with 80-90% lifetime risk of breast cancer) node would be found in upper left of graph as this is a validated SNP variation that has severe consequences (high subjective magnitude), it would be large as cancer risk prevalent in carriers is high, and located to left on X-axis because the behavior for reducing risk, a mastectomy, is hard to perform. A condition with low magnitude that is not easy to do (lower left) are least likely to have successful behavior actions. The action likelihood line, superimposed above, is an approximate behavior change threshold approximated based on Fogg's findings.[14]

The principles of simplicity and clarity was used to limit additional imagery, features such as shape were discussed for the health behavior, but assessment indicated the design would be too busy and confusing. Hue was considered for risk, but judged difficult to interpret.

A small user evaluation was administered using a convenience sample of Survey Monkey Audience members (16 respondents: ages 20-72 yo, mean age 58 yo; 11 female). All respondents were based in the United States. Half had received no formal genetics training while 31% had at least a semester covering genetics topics in high school or college and 18% had greater than a semester of training/education in genetics-related topics. Respondents were questioned regarding the GMV's understandability, usefulness in selecting health behaviors, features, clarity and future enhancements.

DISCUSSION

More than half of respondents found the GMV somewhat understandable and usable, but findings indicate further work is needed on this alpha release. The ability to filter the GMV viewable elements across several data types was the most preferred feature. Improvements should include: clarifying axis label meaning; providing greater contrast when 2 nodes located in same space; and, a need to further simplify. Providing links to external content on relevant SNP information and evidence would likely be useful to users. Future work may include providing features to track one's progress towards risk reduction over time and additional filters for personalization and clarity. DTC-GT offers individuals and health practitioners opportunities to engage around genome-related health improvement opportunities. With further refinement, GMV may be effective in translating multi-dimensional genomic data into meaningful information for health behavior change. Additional work, including experimentation with additional viz elements and modalities, and participatory design with targeted cases, is required to improve the overall understandability and utility of the GMV.

REFERENCES

1] Su, P. (2013). Direct-to-consumer genetic testing: A comprehensive view. *Yale Journal of Biology and Medicine*, *86*(3), 359–365.

2] Bloss, C., Schork, N., & Topol, E. (2011). Effect of Direct-to-Consumer Genomewide Profiling to Assess Disease Risk. *New England Journal of Medicine*, *364*, 524–534.

3] Jackson, L., Goldsmith, L., & Skirton, H. (2014). Guidance for patients considering direct-to-consumer genetic testing and health professionals involved in their care: development of a practical decision tool. *Family Practice*, *31*(3), 341–348.

4] Trepanier, A., Ahrens, M., Mckinnon, W., Peters, J., Stopfer, J., Grumet, S. C., … Farrell, C. (2004). Genetic Cancer Risk Assessment and Counseling: Recommendations of the National Society of Genetic Counselors. *Journal of Genetic Counseling*, *13*(2), 84–114.

5] Hogarth, S., Javitt, G., & Melzer, D. (2008). The current landscape for direct-to-consumer genetic testing: legal, ethical, and policy issues. *Annual Review of Genomics and Human Genetics*, *9*, 161–82.

6] McBride, C. M., Wade, C. H., & Kaphingst, K.A. (2010). Consumers' Views of Direct-to-Consumer Genetic Information. *Annual Review of Genomics and Human Genetics*, *11*(1), 427–446.

7] Yim, S.-H., & Chung, Y.-J. (2014). Reflections on the US FDA's Warning on Direct-to-Consumer Genetic Testing. *Genomics & Informatics, 12*(4), 151–5.

8] Goldsmith, L., Jackson, L., O'Connor, A., & Skirton, H. (2012). Direct-to-consumer genomic testing: systematic review of the literature on user perspectives. *European Journal of Human Genetics: EJHG, 20*(8), 811–6.

9] McGuire, A. L., Diaz, C. M., Wang, T., & Hilsenbeck, S. G. (2009). Social Networkers' Attitudes Toward Direct-to-Consumer Personal Genome Testing. *The American Journal of Bioethics*, *9*(6–7), 3–10.

10] Yach, D. E. A. (2004). Global burden of chronic diseases. Overcoming impediments to prevention and control. *Journal of the American Medical Association*, *291*(21), 2616–2622.

11] McGinnis, J. M., & Foege, W. H. (1993). Actual Causes of Death in the United States. *JAMA: The Journal of the American Medical Association, 270*(18), 2207.

12] Ammerman AS, Lindquist CH, Lohr KN, Hersey J. The efficacy of behavioral interventions to modify dietary fat and fruit and vegetable intake: a review of the evidence. *Prev Med*. 2002;13(1):25–41.

13] Abraham, C., & Michie, S. (2008). A taxonomy of behavior change techniques used in interventions. *Health Psychology*, *27*(3), 379–387.

14] Fogg, B (2009). A behavior model for persuasive design. In *Proceedings of the 4th International Conference on Persuasive Technology* (Persuasive '09). ACM, New York, NY, USA, , Article 40 , 7 pages

[15] Boeldt, D. L., Schork, N. J., Topol, E. J., & Bloss, C. S. (2015). Influence of individual differences in disease perception on consumer response to direct-to-consumer genomic testing. *Clinical Genetics*, *87*(3), 225–232.

[16] Snyder, S., et al. (2015). A Real-World Approach to a Value-Based Decision-Making Framework for Genetic Testing. Journal of Patient-Centered Research and Reviews, 2, 9–10.

17] Nielsen, C. B., Cantor, M., Dubchak, I., Gordon, D., & Wang, T. (2010). Visualizing genomes: techniques and challenges. *Nature Methods*, *7*(3s), S5–S15.

18] Ritchie Lab Software products website. Retrieved October 2016 from https://ritchielab.psu.edu/software.

19] Pendergrass, S. A., Dudek, S. M., Crawford, D. C., & Ritchie, M. D. (2010). Synthesis-View: visualization and interpretation of SNP association results for multi-cohort, multi-phenotype data and meta-analysis. *BioData Mining*, *3*, 10.

20] Nielsen, C. B. (2016). Visualization: A Mind-Machine Interface for Discovery. *Trends in Genetics*, 32(2), 73–75.

Poster #5

Rapid, Iterative Design of reCAP, the REal-Time Care Analysis Platform, to Support Care Redesign

Michael A. Ripperger[1], Colin G. Walsh, MD, MA[1,2,3]
[1]Departments of Biomedical Informatics, [2]Medicine, [3]Psychiatry
Vanderbilt University Medical Center, Nashville, Tennessee

Abstract

Healthcare redesign has relied largely on retrospective, ad-hoc, batched analyses to determine whether clinical process interventions have positively impacted care. These approaches may be 1) slow; 2) hard to scale; 3) not readily reproducible; and 4) may only account for patient-level outcomes and process metrics, not metrics that define whether care redesign is happening as intended. To solve this problem and to support a care redesign effort aimed at optimizing utilization patterns for medically and psychosocially complex patients, a REal-time Care Analysis Platform (reCAP) was rapidly, iteratively designed through multiple design sessions with clinicians and stakeholders in the form of a readily-accessible web application. reCAP was created concurrently with the implementation of a care pilot using modern, open-source software frameworks. The developed dashboard prototype provides user-centered, role-specific visual analyses on filterable, dynamic patient and cohort-level data from electronic health records, audit logs, and financial data. Thematic analysis of initial qualitative evaluation data identify strengths in usability, data manipulation, and supporting information needs. Identified areas for improvement include solidifying clinical relevance and further enabling options for rapid reporting.

Introduction

Preventing unnecessary healthcare services and improving access to care for the underserved remain major goals of contemporary healthcare redesign. For example, achieving optimal care for psychosocially and medically complex patients, thus avoiding preventable readmissions and improving outpatient clinic adherence, has proven to be a persistent challenge for the modern healthcare system[1]. One of the reasons this challenge remains is that restructuring clinical workflows can be difficult and costly. But just as importantly, the analysis of care redesign efforts may rely on slow, batched, often ad-hoc analyses. Moreover, one analytics solution for one care pilot may not scale readily to the next.

An example of active care redesign, the implementation of individualized, longitudinal care plans for complex patients with high utilization of emergency rooms and unplanned inpatient admissions, has successfully reduced inpatient admissions and 30-day readmissions[2]. These interventions are relatively recent and, as such, there is a small but increasing amount of data to support their efficacy. Medical centers across the country are beginning to adopt similar programs, which pair longitudinal care plans with alerts to dedicated teams of hospitalists and extensivists in order to maximize coordination of care each time the most complex patients encounter the healthcare system. Vanderbilt University Medical Center (VUMC), a 1000-bed medical and surgical academic medical center in the Mid-South, has enacted a customized care delivery pilot program for the purpose of optimizing patient care, enhancing continuity, and preventing unnecessary healthcare services.

Because care redesign is expensive, difficult to initiate, and costly to maintain, workflow interventions are often tested via clinical pilots. These pilots may be created and implemented rapidly over weeks to months and, while evaluation metrics are identified to measure their success, there are few solutions to support evaluating them in real-time. One such tool, CoCo, employs a unique blending of visual and statistical analytics for cohort comparisons, but is currently not available for use, requires input of tabular datasets, and must be installed on individual computers[3].

Another challenge in the evaluation of care redesign is the fact that program metrics may be unique to particular care pilots or they may not be tracked in a way that makes them amenable to analytics without a data science or natural language processing. For example, metrics of success may be well-characterized in clinical data, such as process metrics like “reduction in length of stay”; but they may also be conceptual or heuristic such as “improved communication between providers”. During the implementation phase, stakeholders may not have the right data in real-time as to whether or not the program is achieving milestones or even if care delivery is being delivered in adherence with components of the interventions themselves. Solving this problem requires the utilization of multiple sources of hospital data. EHDViz appears to be a promising, open-source, web-based, visual analytics tool for displaying patient data from multiple inputs, but requires significant customization to implement[4].

Finally, quality and clinical care improvement efforts often focus on particular cohorts or groups of patients with common demographics, common diagnoses, or common treatment elements in their care plans. Thus, program impact and changes in patients' care trends must be flexibly aggregatable and filterable by participating stakeholders and clinicians themselves. While solutions such as Tableau exist to provide custom dashboards, it does not allow for the quick development of a lightweight and open source solution and requires instead trained developers to build dashboards in a proprietary format[5].

Visual analytic techniques have the ability to convey the large amounts of data housed within EMRs in an intuitive way, providing a solution for information overload[6]. The deployment of visual analytic tools has been shown to be a promising method for analyzing systemic care processes and adherence to care guidelines[7,8]. Use of color, ability to change visual density, and ability to filter are all important techniques for effectively displaying large amounts of temporal EMR data[9]. For example, EventFlow, succeeding LifeLines and LifeLines2, fulfills these requirements, but lacks support for pre/post analyses[10,11,12]. HARVEST provides hospital encounter timelines for individual patients in real-time through traditional web browsers, but does not enable cohort views or data aggregation[13].

Our objective was to rapidly design and deploy a secure, flexible, usable online visual analytics platform concurrently with a care redesign effort at VUMC. We employed interaction design via multiple design sessions with stakeholders and conducted qualitative evaluation via semi-structured interviews and Think Aloud. Quantitative evaluation is underway and will be described. Our goal was to support a customized, longitudinal care delivery program with a similarly prospective, real-time, readily-accessible visual analytics dashboard that would support rapidly evolving program needs.

Methods

The first step in the development of reCAP at VUMC was a requirements analysis. Three stakeholders were interviewed, representing quality control, pilot leadership, and pilot implementation. A user requirements document was created from the gathered etic data. Core requirements for the application were: real-time display of timelines of patient encounters for thirty initially enrolled patients; ability to filter by patient diagnostic cohorts of interest; time scalability; enabling pre/post intervention comparisons of encounters in aggregate and in subgroups.

Relevant encounters included emergency room encounters, unplanned 30-day readmissions, and outpatient clinic visits. A descriptive workflow model was developed outlining real-world use-cases. The interaction design process was followed throughout the creation of user-interface wireframe mockups with emphasis on understandability and learnability. For the purpose of iterative design, agile principles were followed during the developmental phase and another round of meetings with stakeholders were held after mockup creation. The plan for including audit logs to help analyze whether or not the individualized care plans were being followed was added to the user requirements document. Research on potential existing software solutions did not provide any suitable choice. Focusing on accessibility, the decision was made to develop the application using open-source frameworks. Careful consideration was taken when choosing a secure backend. Django was chosen to provide a RESTful API for the provision of EMR data to the web-browser user interface[14]. User authentication and authorization protocols were followed for security. AngularJS was chosen to create a lightweight, MVC-based frontend[15]. For aesthetics and simplicity, Bootstrap was chosen to stylize the graphic user interface and D3.js was chosen for full customization of graphics[16,17]. Relevant EMR data was accessed through VUMC's research database supported through the Vanderbilt Clinical and Translational Science Award (UL1 RR024975/RR/NCRR, PI: Gordon Bernard). For the protection of patient privacy, Protected Health Information was programmatically de-identified as per HIPAA standards[18]. A stable prototype was developed within one month with the entire developmental timeline spanning three months. A schematic of the design of reCAP is shown (Figure 1).

For the purpose of qualitative evaluation and iterative design, individual, hour-long semi-structured interviews that each included the opportunity for "think-aloud" were conducted. The three relevant stakeholders, being familiar with the novel use-case and prospective end-user requirements, were uniquely suited to be study participants and thus were recruited. During the interview, only audio was recorded, which was then manually transcribed. Two evaluators (MR, CW) reviewed the audio, extracted comments, and conducted thematic analysis[19].

Quantitative evaluation is planned in addition to the qualitative evaluation above. The next iteration of the application will include user audit logs and a click tracking system. Such a logging system will allow the measuring of uptake and usage of specific application features by analyzing user activity trends both collectively and individually.

Figure 1. Design of the developed prototype application and its integration within existing hospital technology. reCAP consists of the 'AngularJS application' and 'Django web server.' .

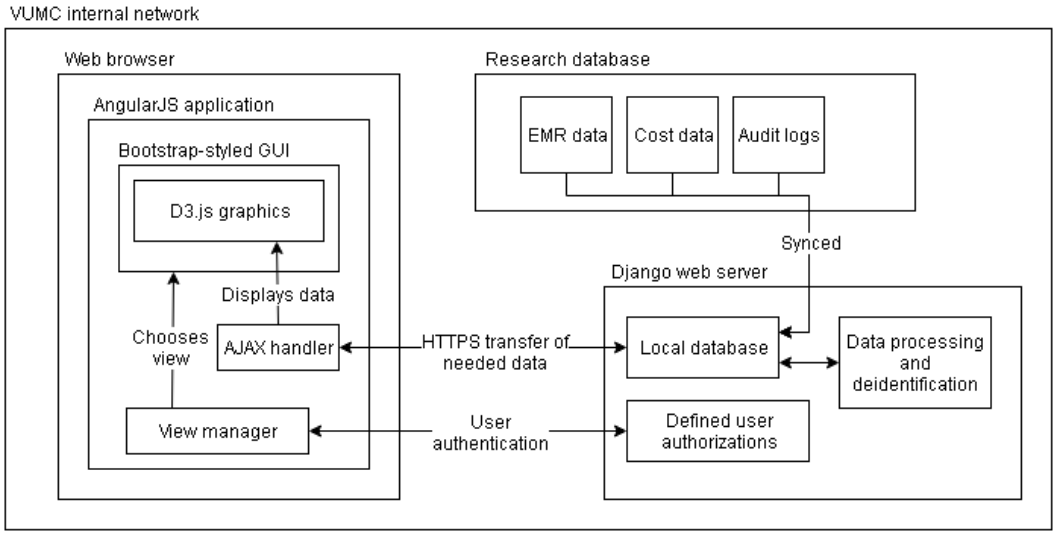

Results

reCAP was completed in accordance with the developed user requirements document (Figure 2). A full outline of the functionality of reCAP is out of scope for this work, but in brief, the core functionality is as follows: timelines of patient encounters give users a comprehensive overview of each patient's clinical trajectory; filtering allows the removal of patient timelines not relevant to the user's purpose or service cohort; aggregating allows the consolidation of descriptive statistics within the panel for each cohort, shifting focus from the individual's trajectory to the group's; and time scaling gives the user the ability to compare encounter histories with the time period of their choosing, whether before or after the care plan start dates.

During the described interviews, study participants interacted with this system using real-world clinical EMR data. Their comments were evaluated via thematic analysis (Table 1).

Figure 2. Screenshot of the main view of reCAP with core user requirements.

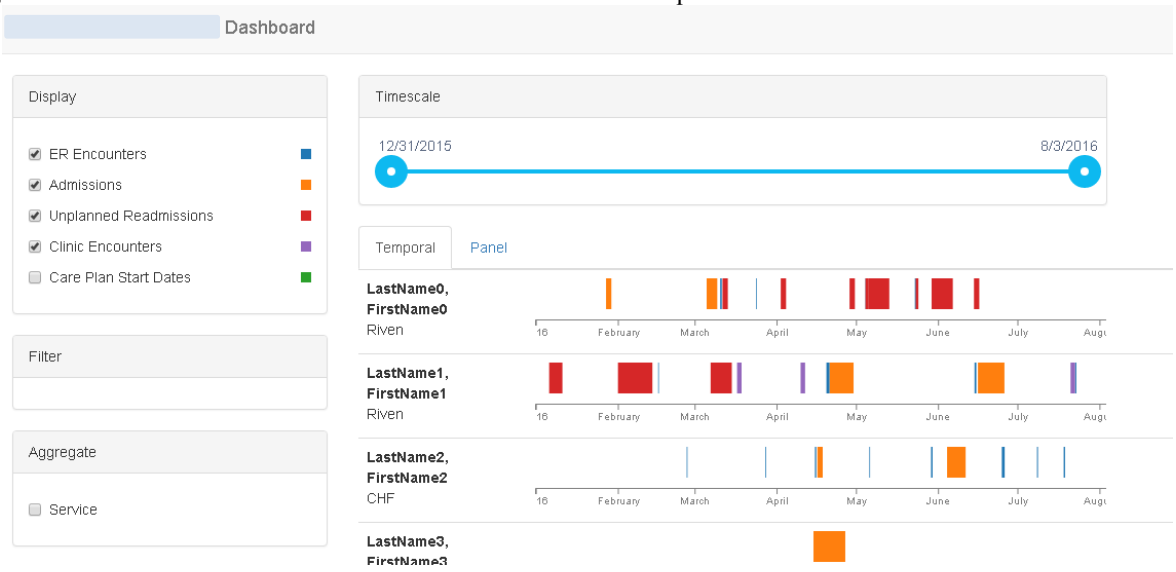

Table 1. Thematic analysis of think-aloud quotes from two participants, representing program quality and program implementation. .

Theme	Participant #1	Participant #2
Accuracy of Displayed Data	"this dashboard is refreshed as of when?" "I wonder to what degree we want to correlate to [other information sources]"	"Ok so I definitely see some of my patients in here" "this must be the encounters that they've had"
Information Needs	"You need to know the numbers behind the percentages" "For me, numerators and denominators are huge" "age and payer source would be valuable additions … zip code might be something else [to add]"	"so we can see when they come to their clinic appointments, that's actually really helpful" "it would be helpful to capture on here whether the plan was actually followed or not" "[adherence to care plans] would be hard to see in real time but that is something I want to know"
Reporting	"and I would add… labeling axes, we will snag it and we will show it. Label it for regular use" "As leaders, we don't want to DO it, we want to minimize data manipulation because we have presentations on a daily or weekly basis"	
Comparison to Existing Solutions	"We're used to snapshotting because every other major view we use is in Tableau" "[Leaders in Quality] are looking at HCAHPs and Press-Ganey. Everyone is looking at those so maybe [they represent visually] a model we can follow"	
Clinical Relevance	"just from a user interface, when I see terms like True/False… your front line doctor would want Yes/No. Even Duration/Time, I'd want "days or hours" to report out" "soften the terminology so it's more clinician friendly" "you might want to put status on here… whether [an encounter] is inpatient or [Observation] Status"	"this is all helpful in terms of data tracking… this is not going to affect workflow for when they come in" "as the inpatient doctor this is not going to matter to me"
Usability	"The nice thing about dashboards like this are they are easy to use, provide easy access and are transparent. Anybody who is involved in this on a daily basis can know the performance of the program and where we might need to continue making changes."	

Discussion

Thematic analysis of qualitative feedback from think-alouds and semi-structured interviews indicated that reCAP had sufficient usability and accessibility but needed improvement in reporting capabilities and further refinement to improve clinical relevance. Subsequent iterative development will further address visual functionality, including color options and change-signifying graphics, thereby improving accessibility and digestibility of data. The hosting of reCAP on a secure internal VUMC server coupled with Django's ability to handle multiple users and their respective rights will provide ample ease of access for every stakeholder. This real-time access will facilitate on-demand control over the customized care delivery program while supporting patient privacy rights. Combining individual patient encounter history timelines with both individual and cohort aggregated analyses of care patterns trends appears to be a powerful tool for allowing insight into program operability and effectiveness. The stark contrast between the visual nature of reCAP and the display of large amounts of data within the EMR user interface upholds that visual analytics is a mechanism of action for improving data comprehension time and user intuition in the healthcare setting.

Timelines became the main focus of reCAP for their ability to visualize individual patient trajectories. Because the care pilot focused specifically on customized care plans at the individual patient-level, its impact was measured at that same level of fidelity. Users consistently emphasized a desire to visualize individual trajectories in a compact, dashboard form. Moreover, because some patients might respond better to this care delivery evolution than others, the potential signal – optimized care patterns for high cost, high need individuals – might have been lost in an aggregation of all utilization for this population, typified by outlier patterns in the first place.

The addition of access timestamps from customized care plan audit logs to the visual timelines, implemented after interviews were conducted, should help stakeholders see whether or not the planned interventions are being followed. Incorporating these logs marks an important feature of reCAP: the ability to utilize multiple data sources not normally available within the EMR user interface. Next, within the developmental pipeline, is the addition of medical center cost data, which will give stakeholders the ability to analyze the scalability of the customized care delivery program as implemented. Integrating this data along with zip codes that can be matched with census data to suggest the socioeconomic backgrounds of enrolled patients will both satisfy user requests and provide a platform for optimizing individualized care at all stages of the program.

Rapid development and the use of modern open-source frameworks has allowed reCAP to be tailored and evaluated in conjunction with the customized care delivery pilot itself. We anticipate that the application may provide robust pre/post analyses of the program's effects on both the hospital and the enrolled patients at any given time, thus empowering stakeholders to optimize and scale with confidence. Currently, reCAP relies upon internally standardized data in order to function. With recent widespread EMR vendor adoptions of SMART-on-FHIR standards through the Argonaut Project, key goals for the future are to support HL7 FHIR standards and to modify the application so it can integrate directly within extant EMRs[20,21]. These future enhancements will permit widespread availability of reCAP and its ability to present visual analyses of the impact of customized care programs prospectively and in real-time.

References

1. Matzer F, Wisiak U V., Graninger M, Sollner W, Stilling HP, Glawischnig-Goschnik M, et al. Biopsychosocial health care needs at the emergency room: challenge of complexity. PLoS One. 2012;7(8).
2. Mercer T, Bae J, Kipnes J, Velazquez M, Thomas S, Setji N. The highest utilizers of care: Individualized care plans to coordinate care, improve healthcare service utilization, and reduce costs at an academic tertiary care center. J Hosp Med. 2015;10(7):419–24.
3. Malik S, Du F, Monroe M, Onukwugha E, Plaisant C, Shneiderman B. Cohort comparison of event sequences with balanced integration of visual analytics and statistics. Proc 20th Int Conf Intell User Interfaces - IUI '15 [Internet]. 2015;38–49. Available from: http://dl.acm.org/citation.cfm?doid=2678025.2701407
4. Badgeley MA, Shameer K, Glicksberg BS, Tomlinson MS, Levin MA, McCormick PJ, et al. EHDViz: clinical dashboard development using open-source technologies. BMJ Open [Internet]. 2016;6(3):e010579. Available from: http://www.ncbi.nlm.nih.gov/pubmed/27013597\nhttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4809078
5. Business Intelligence and Analytics | Tableau Software [Internet]. [cited 2016 Aug 30]. Available from: http://www.tableau.com/

6. Caban JJ, Gotz D. Visual analytics in healthcare – opportunities and research challenges. 2015;(December 2014):260–2.
7. Basole RC, Braunstein ML, Kumar V, Park H, Kahng M, Chau DH orng P, et al. Understanding variations in pediatric asthma care processes in the emergency department using visual analytics. J Am Med Inform Assoc. 2015;22(2):318–23.
8. Basole RC, Park H, Gupta M, Braunstein ML, Chau DH, Thompson M. A visual analytics approach to understanding care process variation and conformance. Proc 2015 Work Vis Anal Healthc [Internet]. 2015;6:1--6:8. Available from: http://doi.acm.org/10.1145/2836034.2836040
9. West VL, Borland D, Hammond WE. Innovative information visualization of electronic health record data: a systematic review. J Am Med Inform Assoc [Internet]. 2015;22(2):330–9. Available from: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4394966&tool=pmcentrez&rendertype=abstract
10. Monroe M, Rongjian L, Hanseung L. Temporal event sequence simplification. IEEE Trans Vis Comput Graph. 2013;(August):2013–9.
11. Plaisant C, Mushlin R, Snyder A, Li J, Heller D, Shneiderman B. LifeLines: using visualization to enhance navigation and analysis of patient records. Proc AMIA Symp [Internet]. 1998;8(98):76–80. Available from: http://www.ncbi.nlm.nih.gov/pubmed/9929185
12. Wang TD, Plaisant C, Shneiderman B, Spring N, Roseman D, Marchand G, et al. Temporal summaries : supporting temporal categorical searching , aggregation and comparison.
13. Hirsch JS, Tanenbaum JS, Lipsky Gorman S, Liu C, Schmitz E, Hashorva D, et al. HARVEST, a longitudinal patient record summarizer. J Am Med Inform Assoc [Internet]. 2015;22(2):263–74. Available from: http://jamia.oxfordjournals.org/content/22/2/263.abstract
14. The Web framework for perfectionists with deadlines | Django [Internet]. [cited 2016 Aug 29]. Available from: https://www.djangoproject.com/
15. AngularJS — Superheroic JavaScript MVW Framework [Internet]. [cited 2016 Aug 30]. Available from: https://angularjs.org/
16. Bootstrap · The world's most popular mobile-first and responsive front-end framework. [Internet]. [cited 2016 Aug 29]. Available from: http://getbootstrap.com/
17. D3.js - Data-Driven Documents [Internet]. [cited 2016 Aug 29]. Available from: https://d3js.org/
18. US Congress. Department of Health and Human Standards for Privacy of Individually. Fed Regist. 2002;1–93.
19. Gale NK, Heath G, Cameron E, Rashid S, Redwood S. Using the framework method for the analysis of qualitative data in multi-disciplinary health research. BMC Med Res Methodol [Internet]. BMC Medical Research Methodology; 2013;13(1):117. Available from: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3848812&tool=pmcentrez&rendertype=abstract\nhttp://www.biomedcentral.com/1471-2288/13/117
20. Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. J Am Med Inform Assoc [Internet]. 2016;1–10. Available from: http://www.ncbi.nlm.nih.gov/pubmed/26911829
21. HL7 Argonaut Project Wiki [Internet]. [cited 2016 Aug 29]. Available from: http://argonautwiki.hl7.org/index.php?title=Main_Page

Acknowledgments

Thank you to the SyBBURE-Searle Undergraduate Research Program at Vanderbilt University, the Department of Biomedical Informatics at Vanderbilt University Medical Center, and the VUMC customized care plan delivery program stakeholders for their continued support of this project.

2016 Workshop on Visual Analytics in Healthcare
Poster #6

Visualizing Patient Genomic Data and Predicted Drug Sensitivities in the SMART Precision Cancer Medicine Application

Krysten Harvey[1], Jeremy L. Warner MD, MS[2-4]
[1]Department of Computer Science, Vanderbilt University; [2]Division of Hematology/Oncology, Vanderbilt University; [3]Department of Biomedical Informatics, Vanderbilt University; [4]Vanderbilt-Ingram Cancer Center, Vanderbilt University Medical Center (all in Nashville, TN, USA)

Abstract: *The SMART Precision Cancer Medicine application compares patient variants of identified cancer genes to frequencies at the population-level. The D3 Javascript library was utilized to encapsulate the results of these comparisons along with suggested drug efficacies within powerful visualizations that are both interactive and intuitive. The aim of this application is to provide support for more personalized decision making at the point-of-care during cancer treatment.*

Background: Precision cancer medicine, i.e., the selection of targeted treatments based on the anticipated efficacy in the context of specific somatic mutations found in a tumor, has proven to be a significant step in the advancement of cancer treatment. Targeting patient care on a genomic level requires tools that provide clinicians a speedy route to access both analyses of patient gene sequencing reports and suggestions for treatment plans tailored to the findings of these reports. Ideally, these functionalities would be consolidated into one application. SMART Precision Cancer Medicine (PCM) is a prototype app[1] that began to explore context-specific gene mutations and variants, using pie and donut charts as the visual vehicles for interaction with PCR-based hotspot mutation tests, which typically detect 0 or 1 mutation in a given tumor specimen. The current work expands on the prototype by using advanced visualizations tying tumor specimen gene variants detected by next-generation sequencing (NGS) panels to expected drug efficacy.

Methods: We enhanced PCM so as to capture higher-dimensional genomic data, as well as to explicitly link genomic variants to efficacy assertions. We explored three types of visualizations to display genomic data: the color-scaled scatter graph, the traditional pie chart, and the dynamic sunburst; the purpose for including all three visualizations within the prototype was to determine the one that clinicians would deem most efficient at the point of care. Each of the three apply different visualization strategies of a proposed information visualization theory.[2] This theory states that initial perception and visual cognition along with the color scheme of an object all contribute greatly to a viewer's interpretation. Furthermore, the scatter-graph uses a two-dimensional graph to scale the frequency of a variant and increasing amino acid position that ensures data is spatially organized in a manner that prevents clustering. On the other hand, the traditional pie chart allows for a clear comparison of frequencies of variants with a single glance. Lastly, the sunburst uses an advanced pattern of interactivity that allows the user to unconsciously develop a cognitive hierarchical structure whose parent branch starts with a view of all the variants of the gene and ends with a single drug. Colors used in each visualization were specifically chosen to decrease the chance of barriers a user could face when interpreting colors, and were derived from the colorblind-safe ColorBrewer[3] scales, with the exception of patient data always being highlighted in yellow and null values being in black. Data from COSMIC[4] and the Jackson Laboratory's Clinical Knowledge Base[5] was used to populate the app. Dynamic links to the external knowledge bases are embedded. Data Driven Documents (D3), a publicly available JavaScript library, was used to implement the visualizations. This library was chosen because it provides an API for manipulating data and displaying results through extensive customization of elements like scalable vector graphics (svgs). The application was developed in the Ionic framework; thus, it is deployable in both iOS and Android operating systems.

Results: The enhanced prototype compares patient variants to population-level variants of the same gene, and provides information specific to each variant (i.e., position, frequency in the population-level dataset, and drug efficacy), as shown in the **Figures**. A mockup of a lung cancer patient with single point mutations in *TP53*, *KRAS*, and *EGFR* (3 mutations, total) was developed for clinical evaluation. The top twenty most frequent mutations found in the COSMIC database for each gene are displayed. The development proceeded with several rounds of input from subject matter experts.

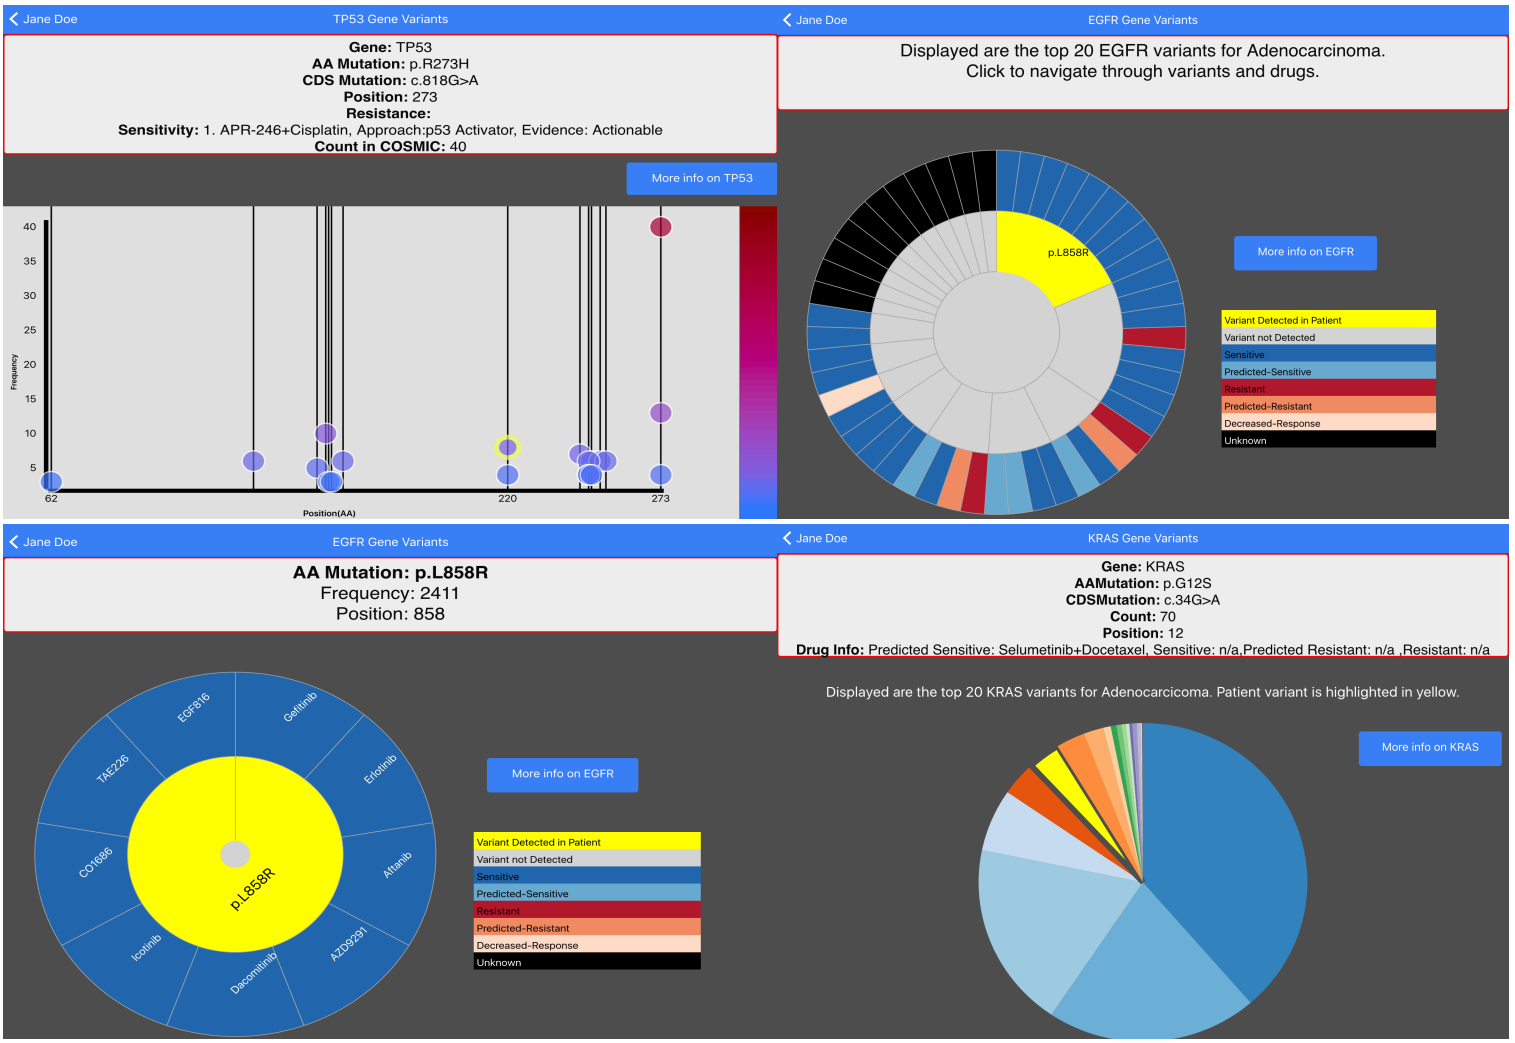

***Figures**:* Screen shots of the app based on a synthetic patient, "Jane Doe." **Upper left:** *TP53* variants, scatter graph; **Upper right:** *EGFR* variants and drug sensitivities, sunburst; **Lower left:** *EGFR*, sunburst with patient variant highlighted, showing drug sensitivities; **Lower right:** *KRAS*, pie chart

Discussion: With an average of 3 variants detected per tumor specimen[6], NGS panels dramatically increase the complexity of cancer care. We anticipate an increasing need for passive clinical decision support (CDS) in the form of information presentation such as described here, as well as active CDS to support treatment decisions. Oncologists and patients are generally not formally trained in data visualization or data science, so it will be important to systematically evaluate various data presentation methods both for clinicians and for patients. In order to test the usability and usefulness of the three visualizations strategies described, we will be introducing this prototype to clinicians through a structured, participant-oriented survey that prompts users to utilize and rate the application's capabilities and provide open feedback. The survey will include comprehensive questions whose focuses range from accessibility and user interface design to how well the visualizations communicate NGS data and drug efficacies.

References:

(1)Warner JL, et al. (2016). "SMART precision cancer medicine: a FHIR-based app to provide genomic information at the point of care." Journal of the American Medical Informatics Association **23**(4): 701-710.
(2)Purchase HC, Natalia A, Jankun-Kelly TJ, Ward M. "Theoretical Foundations of Information Visualization." *InformationVisualization- Human-Centered Issues and Perspectives*.Vol. 4950.Springer Berlin Heidelberg, 2008. 46-64.
(3)Brewer CA. (2013). http://www.ColorBrewer.org, accessed *August 24th, 2016*.
(4)Forbes SA, et al. (2010). "COSMIC (the Catalogue of Somatic Mutations in Cancer): a resource to investigate acquired mutations in human cancer." Nucleic acids research **38**(Database issue): D652-657.
(5)Patterson SE, et al. (2016). "The clinical trial landscape in oncology and connectivity of somatic mutational profiles to targeted therapies." Hum Genomics **10**: 4.
(6)Frampton GM, et al. (2013). "Development and validation of a clinical cancer genomic profiling test based on massively parallel DNA sequencing." Nature biotechnology **31**(11): 1023-1031.

Poster #7

Visualization of EHR Data for Decision-Making in Diabetes and Congestive Heart Failure

Shira Fischer[1], Charles Safran[2], Krzysztof Gajos[3] and Adam Wright[4]

[1]RAND Corporation, [2]BIDMC, [3]Harvard University, [4]Brigham and Women's Hospital

Introduction

The impact of graphical representation of health record data on physician decision-making is important for the design of health information technology.

Methods

Design of a novel simulation environment with data visualizations designed to highlight important clinical trends or relationships, followed by assessment of the impact on decision-making and user experience through a small randomized controlled trial. Providers viewed either a visualization case or a control case and made clinical decisions for patients with chronic diseases.

Results

Twenty-one participants completed the study. The first five results were used to refine the tool; the following sixteen included fifteen primary care physicians and one nurse practitioner, including nine men and seven women. Questions were answered correctly, averaged by individual, 55% of the time. When asked if they noticed objective trends in the data, 85% of the time they said yes. Participants noticed trends more when the visualization was present and found questions more difficult when there was no visualization, but not to statistical significance. No significant variation was found in response correctness by sex, years of experience, or visualization status. Satisfaction with the tool was very high and participants agreed strongly that the tool helped them make better decisions and did so without adding to the time it took them to make these decisions.

Conclusions

The simulation tool allowed us for the first time to test the impact of a visualization on clinician practice in a realistic setting. Designers of electronic health records should consider the ways information presentation can affect decision-making. As trends and relationships can be perceived more easily in graphical format, some lab values and related data may benefit from visual representation. Testing such tools can be done in a clinically realistic context with the right study design. The results of this small study indicate that providers want visualizations and believe that they help them make better and faster decisions.

Poster #8

Visualizing State-Level Chronic Disease Indicators as a Prelude to Insight on Meeting Healthy People 2020 Objectives

Umesh Singh, Ph.D.[1], Victoria Wangia-Anderson, Ph.D.[2],
University of Cincinnati College of Medicine, Cincinnati, OH [1], University of Cincinnati College of Allied Health Sciences, Cincinnati, OH[2];

Abstract

Visualization of the community health status and chronic disease indicators within the most comprehensive and reliable datasets provided by CDC's Division of Population Health provides strategic decision making on future action plans for appropriate healthcare resource allocation, in order to lower disease burden and meet the Healthy People 2020 objectives, an initiative taken by DHHS to identify nationwide health improvement priorities and understand the determinants of health, disease and disability. A user-friendly visualization tool is applied to the community health Status Indicators (CDC) data to generate meaningful insight and development of action plans by concerned officials.

Introduction

Healthy People is an initiative to improve the health of all Americans. Its missions are to identify national health improvement priorities, promote public awareness and understanding of the determinants of health, disease, and disability and the opportunities for progress, to provide measurable objectives and goals that are applicable at the all geographic levels and to involve different sectors to take knowledge and evidence-based steps to improve health policies and care practices.

Chronic diseases account for major causes of death in the United States (1). These include heart disease, cancers, chronic respiratory disease, strokes, diabetes mellitus, chronic kidney disease (2). Smoking, nutritional deficiency or excess, physical inactivity, alcoholism are major risk factors for such diseases (3). Clinical health depends on personal factors such as age, race, gender, dietary/recreational habits. Health of the community, however, is influenced by the prevalence of disease indicators within the community that significantly influence the healthcare burden resulting in high healthcare expenditure for secondary and tertiary care rather than preventive care. Visualization of community health status using relevant indicator helps in prioritizing action plans for improvement in prevalence or incidence rates for diseases that have high healthcare burden. This study was designed to visualize the prevalence of community disease indicators using public datasets Th CHSI dataset retrieved from CDC.gov has key health indicators for local communities and promotes decisive argument about actions plans for improving community health (e.g., obesity, heart disease, cancer). These data are not only useful for public health professionals but also for community members. It has more than 200 measures for the 3,141 US counties. Visualization of the basic epidemiological parameters of these measures can be the primary steps before drilling into the detailed association of relevant indicators Availability of such longitudinal surveillance data on such diseases and risk factors and their appropriate visualization at the different demographic levels (e.g., racial, gender, state, country) helps identify the most vulnerable population groups. Therefore, such visualizations help in planning and implementation of the most effective policies and interventions to resolve public health issues.

Datasets available from Centers for Disease Control and Prevention (CDC) on the chronic disease indicators (CDIs) include a comprehensive set of surveillance indicators that are developed by consensus by the CDC, Council of State and Territorial Epidemiologist (CSTE), and the National Association of Chronic Disease Directors (NACDD). Such datasets enable public health professionals and policymakers to visualize and analyze the well-defined state-level information on chronic diseases and risk factors. These indicators are essential for surveillance, priority listing, and evaluation of public health interventions for chronic disease (4). Sources of CDI data include disease registries, national health surveys, inpatient and emergency department databases, Medicare claims data, policy tracking systems, and the U.S. Census.

Methods

The CDI data file published by CDC in csv format (updated August 23, 2016) was downloaded from the https://chronicdata.cdc.gov/api/views/g4ie-h725/rows.csv?accessType=DOWNLOAD&bom=true. All the variables (e.g., prevalence rates, mortality rate) are adjusted for age (2000 U.S. standard population). After preliminary data cleaning to format length of variable names, the dataset was imported into Tableau v10. A

dashboard is created to effectively demonstrate and compare a specific CDI stratified by gender or race at the state level. The dependent variable named as 'Question' includes the chronic disease indicator of interest, such as 'Asthma prevalence in population aged 18 or above', or 'Hospitalization for asthma' etc. The other variables in the data set named as 'Stratification1' (appearing on the right side of the dashboard) are used for stratification for gender or race. As seen in figure 1A, the 'asthma prevalence above 18 years' as a CDI has been selected from the left panel of the dashboard and compared between whites and African-American population for each state in the two maps at the center of the dashboard. Additional bar-diagram inset within each map lists the top ranking states for each CDI stratified for gender or race. Similarly, asthma prevalence in each state can be compared between male and female populations of the respective states. Using additional dashboards and manipulating the filters two different but related CDIs such as asthma prevalence vs. asthma mortality, or asthma mortality vs. influenza vaccination (%age of population vaccinated) can be compared between states.

Additional dataset, i.e., 'Community Health Status Indicators (CHSI) to Combat Obesity, Heart Disease and Cancer' was retrieved from ftp://ftp.cdc.gov/pub/Health_Statistics/NCHS/Datasets/CHDI/chsi_dataset.zip and the 'DEMOGRAPHICS' file within the zipped file was further visualized in Tableau v10. A disaggregated scatterplot of the 'poverty' data column vs. 'white' column and similar plot between 'poverty' vs. 'black' was visualized. Hovering over the trend line shows statistically significant relationship between these variable. A cluster analysis (new feature in Tableau v10) was then performed for clustering the different States based on these variables as shown in Figure 2, using the k-means clustering algorithm with a variance-based partitioning method that ensures consistency between runs. These clusters can be further drilled down to determine population characteristics within such clusters (Figure 3) and using additional datasets to determine the resource allocation on healthcare, the prevalence of chronic diseases and the associated mortality or modifiable risk factors of specific population groups (e.g., stratified by race) within these states.

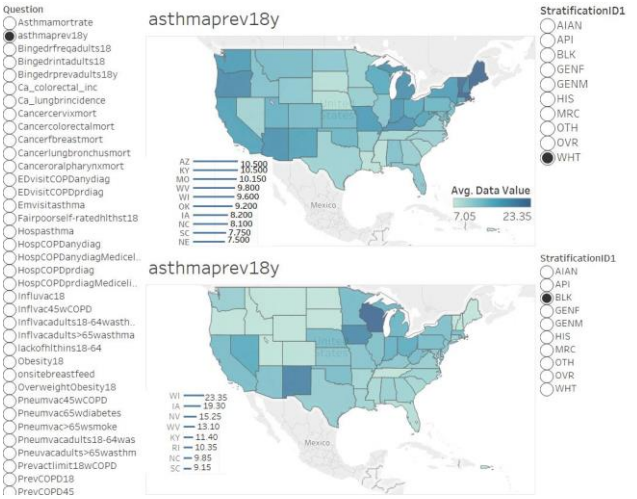

Figure 1. Dashboard to demonstrate chronic disease indicators, stratified by race or gender, at the state level

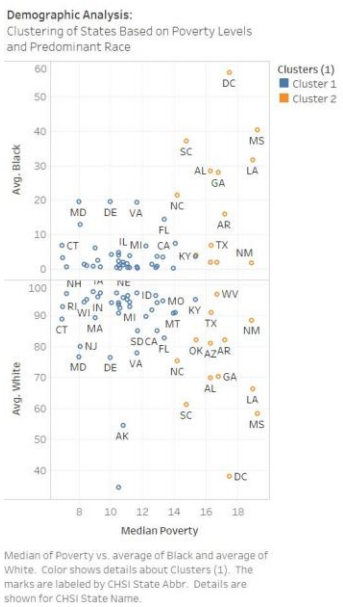

Figure 2. Visualization of State clusters arranged by poverty levels (% population below poverty level, x-axis) and % of population by race on y-axis. States with higher percentage of white population is demonstrated to have lower poverty levels. States within each of the clusters 1 and 2 can be further drilled down to compare healthcare resource allocation, the prevalence of chronic diseases and associated risk factors.

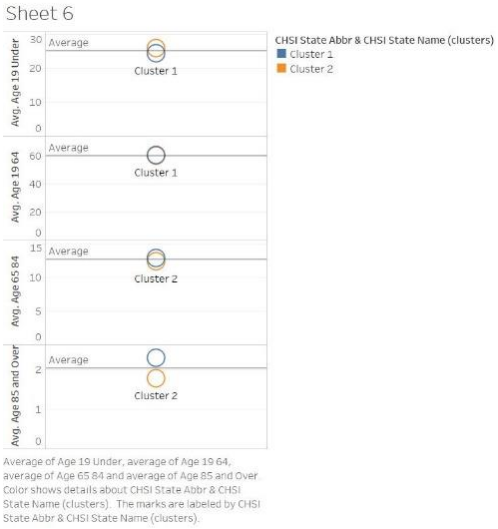

Figure 3. Population age structure in clusters 1 and 2. States with high poverty levels and higher percentage of AA population have distinct population age structures (higher %age of younger dependent population and lower %age of aged) compared to other states.

Conclusion

CDI data is the most appropriate and contemporary collection of chronic disease surveillance data for epidemiologists and other public health officials. Easy and meaningful visualization on such data are possible in Tableau and similar data visualization softwares for comparing the predominant chronic disease prevalence and mortality rates at the state-level stratified by gender and race. These visualizations offer epidemiological intelligence for strategic planning and development of action plans for appropriate allocation of financial and other resources in order to maintain parity on access to health systems and meet the Healthy People 2020 objectives (5).

References

1. Murphy SL, Xu J, Kochanek KD. Deaths: final data for 2010. Natl Vital Stat Rep. 2013;61(4):1-117. Epub 2013/05/08.
2. Heron M. Deaths: leading causes for 2010. Natl Vital Stat Rep. 2013;62(6):1-96. Epub 2013/12/25.
3. Mokdad AH, Marks JS, Stroup DF, Gerberding JL. Actual causes of death in the United States, 2000. JAMA. 2004;291(10):1238-45. Epub 2004/03/11.
4. CDC. Indicators for chronic disease surveillance. MMWR Recomm Rep. 2004;53(RR-11):1-6. Epub 2004/09/11.
5. Cohen R. A brief introduction to the new Healthy People 2020 initiative. Beginnings. 2013;33(1):4-7. Epub 2013/04/05.

LIVE DEMONSTRATIONS

Demo #1:

Evolving Visual Analytics for Better Clinical Decisions

Dave Anderson
Syte Logix, Inc

The exponential growth in digital data to support compound research, new drug development and clinical trials in advancing patient care provides distinct challenges to a clinical researcher. However, this also presents tremendous opportunities for new areas of exploration, cost savings and revenue growth to organizations that are willing to visualize their data in new ways.

Size and complexity of clinical data matters, and unfortunately today's visualization technology does not deliver the critical functionality for the researcher to quickly understand how data is connected and the dependencies between seemingly disparate data sets. Common dashboard visualizations do not provide the necessary context into how data is connected and what insights can be drawn based on these connections.

In order to meet the increasing board pressures to reduce cost and increase ROI, visual analytic tools must evolve to easily support all of the possible data available to researchers, including complex semi-structured, unstructured, and 3rd party data, and enable them to better understand which data is connected and how those data sets are related. This evolution provides the greatest opportunity for companies to use data in a more strategic way to improve value delivered to patients and shareholders.

In this session, we will introduce and provide a software demonstration to show:

- How very large, dense, complex data sets can be quickly and efficiently integrated into a visual analysis program
- A set of visualizations that explore the connections and dependencies across data sets
- A new method to visually analyze data, enabling a deeper, contextual exploration of data
- How customers adopting this new method are realizing tremendous cost savings and improving their competitive position

Demo #2:

A semantic search engine for integration and visualization of Electronic Health Record data and enrichment with open access and public life sciences data sources

Filip Pattyn[1], Bhanu Bah[2] and Hans Constandt[1]
[1]Ontoforce
[2]Harvard Medical School

ONTOFORCE (http://www.ontoforce.com) has developed DISQOVER (http://www.disqover.com), a semantic search engine with faceted search capabilities for life sciences. It currently allows to search automatically across more than 110+ different public data sources that are aggregated, interlinked and contain information about 21 different data types. This system uses semantic web technologies to embrace the mapping efforts from different projects like Unified Medical Language System (UMLS), SNOMED CT, ICD10, ICD9, MedDRA, Human Disease Ontology (DO), Medical Subject Headings (MeSH) and Human Phenotype Ontology (HPO) amongst others. These projects structure and encode information related to diseases, phenotypes, and clinical signs.

Other covered data types are genes, proteins, pathways, drugs, medicines, publications, patents, clinical studies and data types in translational research that are mainly provided via the eagle-i network (http://www.eagle-i.net), such as antibodies, plasmids, biospecimen, cell lines etc.

Eagle-I network at Harvard Catalyst is a linked data open source framework where information about research resources from multiple institutions is collected organized semantically, and shared. It gives any scientist access to one of the largest collections of information about core facilities, stem cells, diagnostic Laboratories, instruments and other valuable, biomedical resources. Institutions collect and share information about research resources in a highly organized and accessible way owing to the Eagle-I Resource Ontology (ERO). Today, over 40 academic and not-for-profit research institutions are represented in the eagle-i open source software community. Continued technical and infrastructure support is funded by Harvard Catalyst at Harvard Medical School (grant number 1UL1 TR001102-01).

Eagle-I project has partnered with Ontoforce to bring internal and public data closer to the researcher, greatly reducing the workload of laborious searching and manual data aggregation from many different sources. Through a semantic data curation platform, such as eagle-i, and semantic search platforms, such as DISQOVER™, a researcher only needs to access a single, user friendly and consistent platform. The

navigable visualization turns searching into an intuitive experience. The different internal databases have already been connected to the semantic search platform, and all data has been categorized through ontologies. With linked data technology, logical but often hidden connections between various data points are brought to light leading to better insights on patient care.

Electronic Health Record (EHR) data often lacks structure and encoding schemes. DISQOVER has the capability to deal with these less structured information sources and can apply the knowledge from public data and generally accepted ontologies or classification schemes to bring more structure and links. An example are the capabilities to combine the usage of different disease classifications with the structuring of clinical signs via HPO. This means that diseases or conditions encoded as SNOMED CT, ICD10 or other classification codes can be linked to HPO phenotypes and internally linked to individuals behind the organization's firewall but also enriched with links to literature, pathways and other data sources. A medical practitioner can be enabled to access the patient specific data, summarize and visualize this data by disease, disease group, phenotype or other parameters and can compare this with information from other patients. During the demo we will showcase some real life examples in this space merging public datasets but also showcasing how internal land/or tranSMART sources are easily integrated and get ready to be used broadly.

By default, DISQOVER is an open access search engine with a consistent interface for searching in publicly available data sources and is coined 'public DISQOVER'. An automated data update process is running to be as much as possible in sync with the original sources. At every moment it is possible to show the origin of every individual data point to guarantee complete openness in data provenance. The 'public DISQOVER' system can be interpreted as a semantic web enriched mirror of all comprised data sources. This is necessary to generate a fast and interactive web application, which is currently not possible with a classical semantic web data federation approach due to the technical issues with responsiveness and the lack of data sources available as accessible semantic data endpoints.

A second DISQOVER type is the 'internal DISQOVER' setup. The search engine can be installed inside an organization with additional features like user group management for data access and data visualization. An 'internal DISQOVER' allows to structure, integrate, link and visualize internal data from an organization. The system keeps the information completely secure from the outside world. Moreover, an 'internal DISQOVER' can communicate with the central 'public DISQOVER' using a specifically developed 'DISQOVER data federation' system to enrich this internal data with the wealth of already structured public data. With this approach, no internal information is transferred to the pubic DISQOVER. There is only a unidirectional data flow from the public system to the internal DISQOVER.

The system has the flexibility and can be configured to integrate different data sources. Configuration scripts give you full control how to integrate the data, what

and how to show the data for specific user roles. The data can be easily filtered via customizable filters or facets. Specific facet visualizations are available for different types of information: geographical maps for locations, tree view for hierarchical data, timelines for dates etc.

The primary goal is to allow every user to find what she/he needs in a way that a user understands best, via his or her/his understanding of the knowledge domain and without having to know the exacts search terms or labels. And results even link out t interesting (research or health-economics) information outside of the user's expertise. DISQOVER is a platform usable for all stakeholders in biomedical research and healthcare with a strong focus on the user experience and the user interface where we strive to make 'everybody a data scientist'.

Demo #3

A Clinical Data Analytics Workbench to Streamline Analytics Tasks and Visualize Key Results

Yiqin Yu, MS[1], Xiang Li, PhD[1], Haifeng Liu, PhD[1], Bibo Hao, MS[1], Wen Sun, PhD[1], Guotong Xie, PhD[1]
[1]IBM Research, Beijing, China

Abstract

Recent advances in big data and machine learning related technologies provide powerful capabilities for clinical researchers to gain insights from massive healthcare data. However, compared to traditional clinical research methods, what kind of analytical results are new, and most importantly, with meaningful clinical implications, is still unclear. This demo introduces the Clinical Data Analytics Workbench, which streamlines different kinds of clinical data analytics tasks, and visualizes the key results to provide clinical insights. With this workbench, clinical researchers could leverage the power of modern (big) data analysis capabilities, and enhance the implementation of their clinical research tasks.

Introduction

Recent advances in big data and machine learning related technologies provide powerful capabilities for clinical researchers to gain insights from massive healthcare data [1]. Comparing to performing clinical data analysis in conventional software or tools (e.g., SPSS, SAS), developing with Python, R, etc. enables researchers to take advantage of big data analytics tools (e.g., Apache Spark) and have more flexible control over the analytics process. However, comparing to traditional clinical research methods, clinical researchers still have questions. First, what kind of new results can be generated from modern technologies? Second, do these new results contain meaningful and explicable clinical implications? This demo introduces a web-based application named Clinical Data Analytics Workbench, which design and streamlines different kinds of clinical data analytics tasks, and visualizes the key results to provide clinical insights. With this workbench, clinical researchers can leverage the power of modern (big) data analysis capabilities, and enhance the implementation of their clinical research tasks.

Streamlining of Clinical Data Analytics Tasks

The design of different types of clinical data analytics tasks are based on the discussion between clinical researchers in hospitals and data scientists who have computer science and statistics background. Focusing on structured clinical data such as Electronic Medical Record and Registry data, there are a set of clinical research tasks been identified, including risk prediction, patient stratification, treatment effectiveness analysis, and clinical pathway analysis. The different types of tasks are initially raised by clinical researchers and further confirmed by the results of data exploration. For example, after exploring the distribution of different disease events, certain disease is selected as the target outcome of a risk prediction task. In this demo we will focusing on the first two tasks. The dataset used in the demo contains registry data for about 17,000 patients, who were in treatment in the department of cardiology of a hospital in Beijing, China from Aug, 2011 to Jun, 2015.

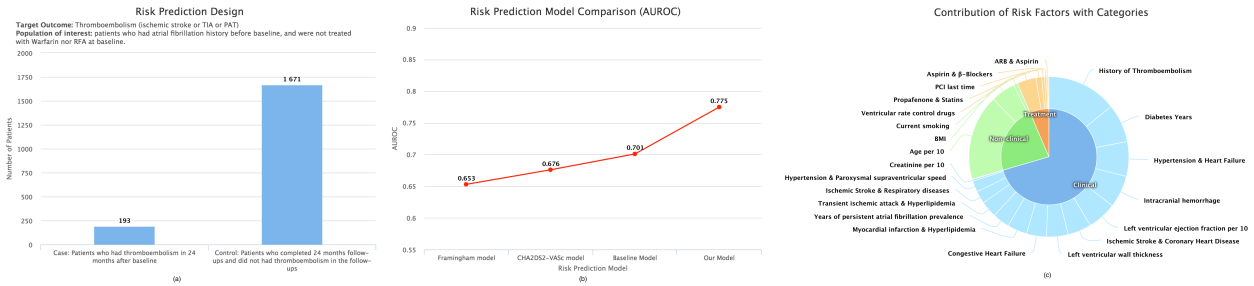

Figure 1. Key results visualization in risk prediction task. (a) Design of the risk prediction task. (b) model comparison on AUROC. (c) Exploration of contribution of factors in categories.

Visualization of Risk Prediction Task

The Risk Prediction task aims to build a sound predictive model between the target outcome and related features. The task is streamlined into four steps: 1) **Task Design**. Task Design (Figure 1(a)) specifies the target outcome of the prediction task, the population of interest, and interested features which should be used for prediction. 2)

Predictive Modeling. This step builds risk models based on different prediction/classification algorithms and compares the performance of these models (Figure 1(b)). The most commonly used measurement is the area under the receiver operating characteristics curve (AUROC). 3) **Factor Exploration**. The results of this step are critical for clinical researchers, as the ultimate goal of a risk prediction task is to find factors with clinical implications for the target outcome (Figure 1(c)). Traditional risk prediction tasks usually have limited number of factors and report the contribution of them in a pie chat. By using modern technologies, theoretically unlimited number of factors can be analyzed and involved in the risk model. Here we use a donut chart with two layers to represent the factors. The inner layer summarizes the distribution of contribution in categories such as clinical, non-clinical and treatment, which provides a high level view about what kind of factors impact the outcome mostly. And the outer layer shows the contribution for each factor. 4) **Risk Distribution**. This step calculates the risk score for each case (patient) in the dataset with the selected risk model, and shows the distribution in different risk levels.

Visualization of Patient Stratification Task

The Patient Stratification task tries to stratify patients in a cohort into multiple groups by cluster analysis algorithms, which might have different health conditions and will lead to different treatment options. Steps include: 1) **Cluster Projection**. After the population of interest is defined, cluster analyses are run and each case is assigned to certain group (cluster). All cases are projected to a 3-dimention space and visualized in a 3-D scatter chart as well as their group label. 2) **Group Characteristics**. The characteristics of each group contain two main information (Figure 2): 1) the distribution of all features used in the clustering algorithm, which is illustrated in a spider chart, and 2) the leading features that have a particularly high distribution on the group (red colored features), or vice versa (green colored features). 3) **Grouping Rules Mining**. In a real hospital environment, physicians usually perform patient stratification by rules. For this purpose, grouping rules are mined by a decision tree according to the clustering results (Figure 3(a)), which can be treated as a provenance of how the rules are generated by involving important features into the decision tree one by one. The integrated rules are listed by groups as well as the patient number and confidence score (Figure 3(b)).

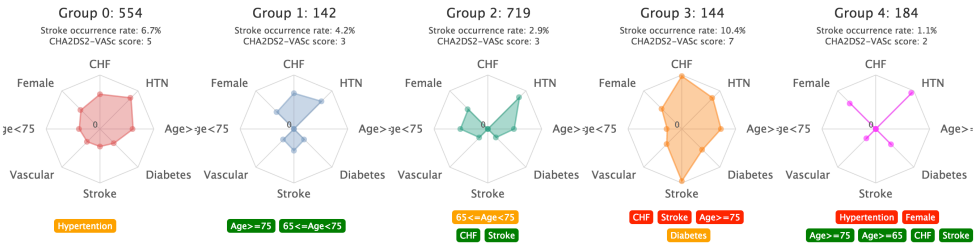

Figure 2. Group characteristics visualization in patient stratification task.

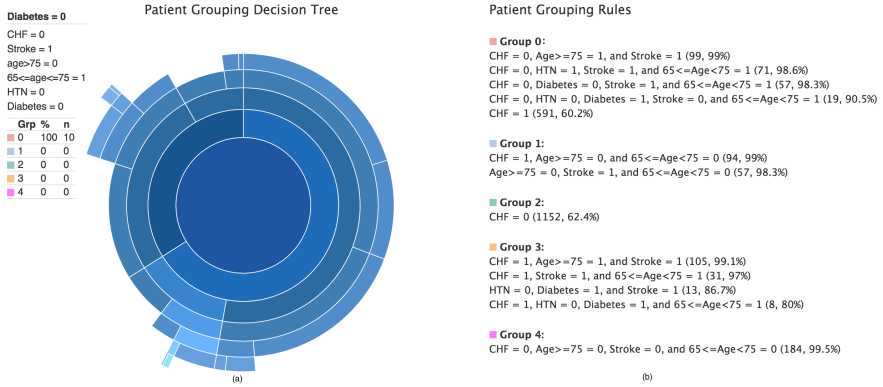

Figure 3. Group characteristics visualization in patient stratification task. (a) The decision tree for grouping patients according to the cluster analysis results. (b) Summarized grouping rules derived from the decision tree.

References

1. Bibo HAO, Wen SUN, Yiqin YU, Jing LI, Gang HU, Guotong XIE. Accelerate Healthcare Data Analytics: An Agile Practice to Perform Collaborative and Reproducible Analyses. MIE 2016.

Demo #4

Demo: Peeking into Patterns of Clinical Event Sequences with Peekquence

Bum Chul Kwon[1], **Janu Verma**[1], **Adam Perer**[1]
[1] **IBM T.J. Watson Research Center, Yorktown Heights, NY, USA**

1 Introduction

Finding temporal patterns in longitudinal event sequences is a challenging task, as the volume and variety of events often make it difficult to extract salient patterns. In response to this challenge, data scientists have turned to machine learning, known as frequent sequence mining (FSM) techniques, to automatically detect the most common sequences of events to unearth interesting patterns. For instance, Frequence [4], Care Pathway Explorer [5], and TimeStitch [6] all use frequent sequence mining techniques to find frequent sequences of events.

However, these algorithms often require users to specify a support threshold that, if too high, will yield only a few patterns, or if too low, will yield numerous patterns that may be difficult for data scientists to determine the interesting sequences from the mundane. In this demo, we aim to make the results of frequent sequence mining algorithms more interpretable by giving end-users powerful ways to explore the data.

Our novel visual analytics system, Peekquence [1], integrates several new techniques that include: 1) powerful ways to navigate the patterns by sorting with metrics relevant to users (variability, correlation to outcome, etc), 2) integration of patterns with patient timelines, so users can understand where the patterns occur in the actual data, and 3) overviews the summarize the most common events in the patterns.

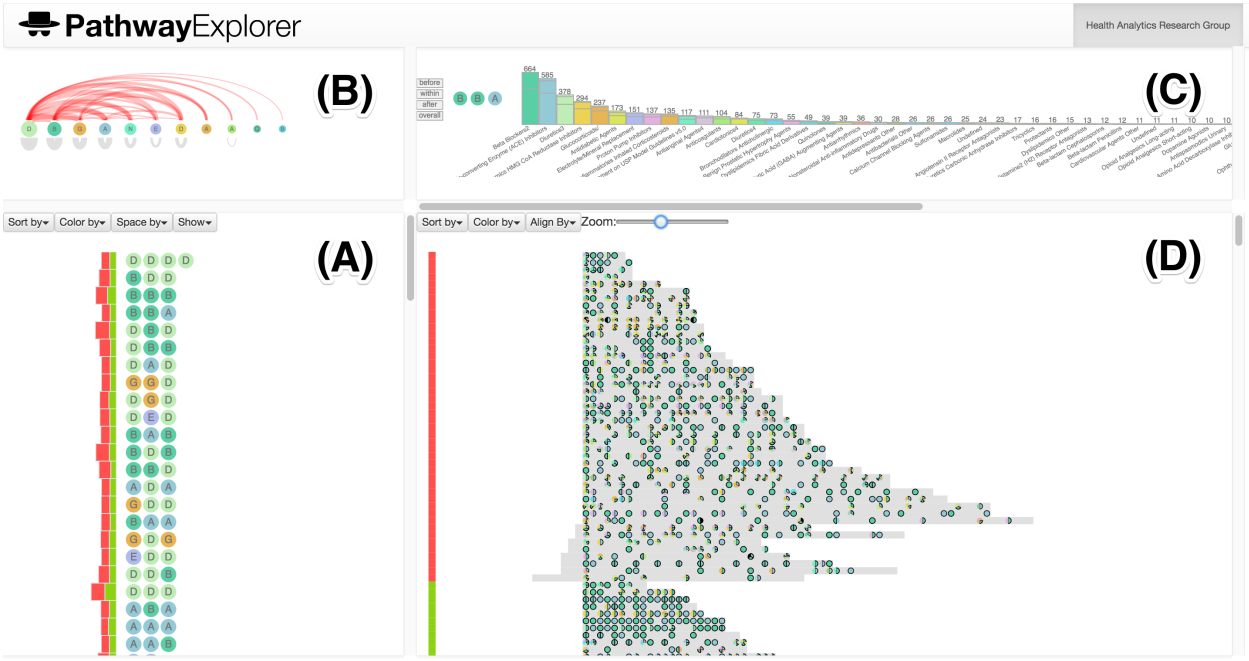

Figure 1: Peekquence consists of four views: (A) the pattern list view showing patterns mined from SPAM with event sequences (colored circles with letters) as well as bars of patients with the ratio of case and control labels (diagnosis of a disease); (B) the sequence network view showing the frequency of event type co-occurrences within mined patterns; (C) the event co-occurrence histogram view showing the frequency of events co-occurring for a selected pattern; (D) the patient timeline view showing patients' event sequences that match the selected pattern.

2 Peekquence

The core visual unit of Peekquence are mined patterns, rather than events. As patterns may contain many different event types and be composed of long event sequences, visualization techniques based on sankey diagrams (a la CareFlow [3]) or aggregated vertical bars (a la EventFlow [2]) tend to suffer from visual complexity without user-controlled filters based on domain expertise. Instead, we opted for a simpler visualization technique: a list of patterns, made up of *event glyphs* that visually encode each event type in the pattern. The event glyphs are visually encoded as circles, colored according to an categorical ontology, and labeled with an abbreviation of the event type's name. All of the four views in Peekquence, shown in Figure 1, use this glyph as the common visual element. In addition to a list of patterns (Figure 1A), there is an overview of common event types in the patterns (Figure 1B), histograms that summarize event types that co-occur with the patterns (Figure 1C), and a coordinated view to the actual patient timelines to understand how the mined patterns manifest in the actual data (Figure 1D).

In Figure 1, Peekquence is demonstrated on mined patterns from a cohort of patients with diagnoses of both congestive heart failure (CHF) and chronic obstructive pulmonary disease (COPD). Of these patients, some are cases that were hospitalized and the remaining are matched controls who have the disease but were not hospitalized. The goal is to use Peekquence to understand if the patients with CHF and COPD that were hospitalized have any distinct patterns of treatments compared to those who were not hospitalizd. For these patients, one year of data is mined after their diagnoses of CHF. In this figure, only treatment events are mined, but the system is capable of merging multiple types of events (e.g. diagnoses, procedures, and labs).

Peekquence has led to interesting discoveries of the benefits and problems by relying on mined patterns as the main unit of visualization. There was no data curation done to the event types loaded into the user interface, but the algorithm was able to surface highly relevant types due to their prominence among patients with CHF and COPD.

3 Conclusion

In this paper, we presented Peekquence, a visual analytics system which aims to increase the interpretability of frequent sequence mining algorithms. The four views combined with interactions provide useful functionalities for users to make sense of patterns as well as their occurrences within patients' records.

References

[1] B. C. Kwon, J. Verma, and A. Perer. Peekquence: Visual analytics for event sequence data. *ACM SIGKDD Workshop on Interactive Data Exploration and Analytics (IDEA 2016)*, 2016.

[2] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman. Temporal Event Sequence Simplification. *IEEE Transactions on Visualization and Computer Graphics*, 19, 2013.

[3] A. Perer and D. Gotz. Data-driven Exploration of Care Plans for Patients. In *ACM Human Factors in Computing Systems*, pages 439–444, 2013.

[4] A. Perer and F. Wang. Frequence: Interactive mining and visualization of temporal frequent event sequences. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*, pages 153–162, New York, NY, USA, 2014. ACM.

[5] A. Perer, F. Wang, and J. Hu. Mining and exploring care pathways from electronic medical records with visual analytics. *Journal of Biomedical Informatics*, 56(C):369–378, Aug. 2015.

[6] P. J. Polack Jr, S.-T. Chen, M. Kahng, M. Sharmin, and D. H. Chau. Timestitch: Interactive multi-focus cohort discovery and comparison. In *IEEE Proceedings of the Visual Analytics Science and Technology (VAST)*, pages 209–210. IEEE, 2015.

Demo #5

Using Advanced Analytics and Visualization to Detect HIPAA Violations

Nick Culbertson and Robert Lord
Protenus, Inc

Abstract:
Patient privacy violations in the EHR are all-too-common occurrences, but accurately detecting them is a huge and thus far nearly-insurmountable challenge. Examples abound of both everyday and truly horrific violations of privacy that occur in EHRs, ranging from co-workers snooping on each other, to criminal networks systematically mining patient identity data through bribing clinical staff. Detecting these inappropriate accesses to patient data is a huge challenge, made even more difficult by the current dearth of robust, big-data solutions driven by modern technology and contemporary visualization tools.

The Protenus platform delivers a new approach to the patient privacy monitoring challenge. Critically and uniquely, Protenus incorporates an understanding of the complex clinical environment, as well as next-generation visualizations that demonstrate the difference between appropriate and aberrant behavior. This advanced approach to privacy fosters a renewed sense of organizational cooperation, decreasing undue stresses on compliance and security teams and replacing traditional compliance technology. Protenus streamlines threat detection and resolution cycles from months to a matter of minutes.

The authors seek to demonstrate their platform's methods and visualizations to provide a different perspective on the ability of clinical informatics and health data visualization to play a whole new role in privacy and security.

Demo #6

Visualizing State-Level Chronic Disease Indicators as a Prelude to Insight on Meeting Healthy People 2020 Objectives

Umesh Singh[1], Victoria Wangia-Anderson[2],
[1]University of Cincinnati College of Medicine, Cincinnati, OH
[2]University of Cincinnati College of Allied Health Sciences, Cincinnati, OH

Abstract:
Visualization of the community health status and chronic disease indicators within the most comprehensive and reliable datasets provided by CDC's Division of Population Health provides strategic decision making on future action plans for appropriate healthcare resource allocation, in order to lower disease burden and meet the Healthy People 2020 objectives, an initiative taken by DHHS to identify nationwide health improvement priorities and understand the determinants of health, disease and disability. A user-friendly visualization tool is applied to the community health Status Indicators (CDC) data to generate meaningful insight and development of action plans by concerned officials.

Demo #7:

Visualizing Patient Genomic Data and Predicted Drug Sensitivities in the SMART Precision Cancer Medicine Application

Krysten Harvey[1] and Jeremy L. Warner[2,3,4]

[1]Department of Computer Science, Vanderbilt University;
[2]Division of Hematology/Oncology, Vanderbilt University;
[3]Department of Biomedical Informatics, Vanderbilt University;
[4]Vanderbilt-Ingram Cancer Center, Vanderbilt University Medical Center

Abstract:
The SMART Precision Cancer Medicine application compares patient variants of identified cancer genes to frequencies at the population-level. The D3 Javascript library was utilized to encapsulate the results of these comparisons along with suggested drug efficacies within powerful visualizations that are both interactive and intuitive. The aim of this application is to provide support for more personalized decision making at the point-of-care during cancer treatment.

Demo #8:

Genetics-based Motivation Viz ("GMV"): Visualizing Direct-to-Consumer Genetic Test Results To Empower Health Behavior Change

Kenyon Crowley[1,2,3,4], Joohee Choi[1], Rohan Bondili[1]

[1]University of Maryland iSchool, College Park, MD;
[2]Human-Computer Interaction Lab (HCIL), College Park, MD;
[3]Robert H. Smith School of Business, College Park, MD;
[4]Center for Health Information & Decision Systems (CHIDS), College Park, MD

Abstract

The human genome is a complex mix of approximately 10 million SNPs (single nucleotide polymorphisms). Each SNP represents a difference in a single DNA building block and may be correlated with traits and health conditions. The translation of complex multidimensional genetic data into actionable health behavior information to prevent disease is a difficult task. Our team developed the Genetics-based Motivation Viz ("GMV") in order to help individuals and health practitioners identify, motivate and prioritize genetics-based health behavior change opportunities. GMV incorporates an abstraction of a validated health behavior change framework, in combination with an individual's direct-toconsumer genetic testing (DTC-GT) results (including their health condition risk, moderating risk factors, strength of evidence supporting claim), juxtaposed with their self-reported ease of doing certain health behaviors (such as eating a low fat diet). GMV was deployed in a web browser and evaluated via a survey with 16 respondents. The majority of respondents indicated GMV was at least somewhat useful and understandable, but also identified areas for future work. With further development, GMV may be effective in translating genomic data into meaningful health behavior change information useful for individuals and health practitioners.

Demo #9:

Rapid, Iterative Design of reCAP, the REal-Time Care Analysis Platform, to Support Care Redesign

Michael A. Ripperger and Colin G. Walsh
Vanderbilt University Medical Center, Nashville, Tennessee

Abstract:
Healthcare redesign has relied largely on retrospective, ad-hoc, batched analyses to determine whether clinical process interventions have positively impacted care. These approaches may be 1) slow; 2) hard to scale; 3) not readily reproducible; and 4) may only account for patient-level outcomes and process metrics, not metrics that define whether care redesign is happening as intended. To solve this problem and to support a care redesign effort aimed at optimizing utilization patterns for medically and psychosocially complex patients, a REal-time Care Analysis Platform (reCAP) was rapidly, iteratively designed through multiple design sessions with clinicians and stakeholders in the form of a readily-accessible web application. reCAP was created concurrently with the implementation of a care pilot using modern, open-source software frameworks. The developed dashboard prototype provides user-centered, role-specific visual analyses on filterable, dynamic patient and cohort-level data from electronic health records, audit logs, and financial data. Thematic analysis of initial qualitative evaluation data identify strengths in usability, data manipulation, and supporting information needs.

Demo #10:

Visual Summarization of Longitudinal Clinical Trajectories to Improve Population Health Analysis

Filip Dabek and Jesus J Caban

Walter Reed National Military Medical Center
Department of Defense

Abstract:
Summarizing a collection of temporal sequences is a difficult task given the irregular and variable patterns often found in longitudinal events. Across a wide array of domains, researchers and analysts seek to determine ways to identify common temporal paths, to build trajectories between individual events, and to understand the relationships between different events. While these tasks continue to be difficult on small and structured datasets, they are increased tenfold on temporal sequences that are noisy, irregular, and voluminous in size. Approaches that enable analysis of temporal sequences of large or small, noisy or clean, irregular or structured datasets open new opportunities to identify key information embedded within longitudinal data. We demonstrate a scalable framework that has been designed to visually explore large collections of temporal sequences by combining advanced event mining algorithms with visualization techniques to overcome some of the challenges and complexities of the data. The system has been tested with a comprehensive clinical dataset of 98,342 patients and 8.7 million longitudinal events showing the effectiveness of the techniques within large and complex datasets.

Program Committee

Organizers:

Jesus J Caban, PhD	Walter Reed National Military Medical Center
Bhanu Bahl, PhD	Harvard Medical School
Jacquelyn W. Blaz, PhD, MS	UW-Madison School of Nursing
Sary Beidas, MD	Orange Park Medical Center
Carolyn Petersen, MBI, MS	Mayo Clinic
Danny Wu, PhD, MSI	University of Cincinnati
Adam Perer, PhD	IBM Research
Jeremy Warner, MD	Vanderbilt University
Uba Backonja, PhD RN	University of Washington School of Medicine
Shira Fischer, MD, PhD	RAND Corporation