Composer—Visual Cohort Analysis of Patient Outcomes

Jen Rogers¹ Nicholas Spina² Ashley Neese² Rachel Hess³ Darrel Brodke² Alexander Lex¹

¹Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, Utah, United States

²Department of Orthopedics, University of Utah, Salt Lake City, Utah, United States

³Department of Population Health Sciences, University of Utah, Salt Lake City, Utah, United States

Address for correspondence Alexander Lex, PhD, Scientific Computing and Imaging Institute, University of Utah, 72 Central Campus, Dr Salt Lake City, UT 84112, United States (e-mail: alex@sci.utah.edu).

Appl Clin Inform 2019;10:278–285.

Abstract

Objective Visual cohort analysis utilizing electronic health record data has become an important tool in clinical assessment of patient outcomes. In this article, we introduce Composer, a visual analysis tool for orthopedic surgeons to compare changes in physical functions of a patient cohort following various spinal procedures. The goal of our project is to help researchers analyze outcomes of procedures and facilitate informed decision-making about treatment options between patient and clinician.

Methods In collaboration with orthopedic surgeons and researchers, we defined domain-specific user requirements to inform the design. We developed the tool in an iterative process with our collaborators to develop and refine functionality. With Composer, analysts can dynamically define a patient cohort using demographic information, clinical parameters, and events in patient medical histories and then analyze patient-reported outcome scores for the cohort over time, as well as compare it to other cohorts. Using Composer's current iteration, we provide a usage scenario for use of the tool in a clinical setting.

Conclusion We have developed a prototype cohort analysis tool to help clinicians assess patient treatment options by analyzing prior cases with similar characteristics. Although Composer was designed using patient data specific to orthopedic research, we believe the tool is generalizable to other healthcare domains. A long-term goal for Composer is to develop the application into a shared decision-making tool that allows translation of comparison and analysis from a clinician-facing interface into visual representations to communicate treatment options to patients.

Keywords

- ▶ cohort analysis
- ▶ visualization
- ▶ support
- ▶ comparisons

Background and Significance

Determining the best treatment option for patients with back pain involves an assessment of their medical histories and a comparison to similar patients. Such comparisons have relied on a physician's memory of related prior cases, which can be influenced by cognitive biases. With an increasing amount of data available for patient populations in electronic health records (EHRs), visual cohort analysis has gained attention as an informative analytic tool in healthcare. Recent work has shown the efficacy of using subsets of similar patients, referred to as cohorts, for outcome analysis and prediction in a “patient-

like-me” approach.^{1,2} This approach can help clinicians assess treatment options for patients with certain characteristics or preexisting conditions (comorbidities) that can influence recovery and response to treatment.

In this article, we introduce Composer, a visual analysis tool for comparison of patient outcomes in cohorts under alternative treatment options. Composer was developed in collaboration with domain experts at the University of Utah Orthopedic Research Center. We incorporate outcome scores that are frequently measured over the course of treatment in the decision-making process, supplementing physicians' memory of prior cases. We used the Patient-Reported Outcomes

received

January 2, 2019

accepted after revision

March 11, 2019

© 2019 Georg Thieme Verlag KG
Stuttgart · New York

DOI <https://doi.org/10.1055/s-0039-1687862>.
ISSN 1869-0327.

Measurement Information System (PROMIS)³ scores as the metric for patient physical function (PF) and well-being over time.

The technical contributions of Composer include methods to flexibly define multiple patient cohorts based on EHR data and demographic attributes as well as medical codes associated with a given medical visit. We provide functionality for PROMIS score normalization to allow for alignment of score trajectories based on events in patient medical histories, such as surgery or injection. We also provide the ability to normalize scores from absolute measurements to relative change to identify improvement of patient PF. Finally, we introduce aggregation methods to deal with larger patient cohorts.

Background

Cohort Analysis

Most clinical guidelines are based on evidence from clinical trials and controlled studies. However, data collected from clinical trials, often sourced from a general population, may not provide an accurate reflection of potential outcomes for subsets of patients with preexisting conditions and comorbidities.² Clinicians are, therefore, interested in using EHR data and observational studies to better identify factors that can influence the recovery of such patients.⁴ A cohort is defined as a subset of the general population that shares one or more defining characteristics. The analysis of cohorts has proven effective in the medical community for identifying factors that affect patient recovery and treatment.

In clinical applications, cohorts can be defined by utilizing patient data collected through the EHR system. The medical community has relied on cohort subsets sourced from a large body of EHR data that can be used for retrospective analysis.^{4,5} Cohorts of patients formed from EHR data have the potential to be used for “patients-like-me” comparisons,² in which clinicians can define a cohort with attributes mirroring a given target patient. These comparisons can help identify factors that influence patient recovery and have been used to develop predictive tools that help domain experts determine the best treatment options for a given patient.^{6–8}

PROMIS Score System

PROMIS is a validated measurement system that evaluates a range of patient PFs.⁹ In this article, we use only PROMIS PF scores. The PROMIS system defines the abilities of a patient with a specific score, which is determined by patient response to a series of questions.³ A patient who can run 10 miles without difficulty would have a PROMIS PF score of approximately 72, whereas a patient with a score of 32 can stand for a short period of time without difficulty.¹⁰ If patients have answered that they have trouble walking a mile, later questions will focus on a smaller range of physical abilities. The score system is converted to a t-score metric that ranges from 0 to 100, with an average ability score of 50 and a standard deviation of 10. All scores are scaled to values relative to the average score. For example, a score of 40 implies PF that is one standard deviation lower than the score of the reference mean.³

The University of Utah Orthopedic Research Center has been a proponent in the use of PROMIS scores to assess patient outcomes.¹¹ Recent research into PROMIS PF scores to evaluate a given procedure relative to cost has identified PROMIS PF as a more accurate assessment of physical well-being for patients with spinal ailments than the Oswestry Disability Index, which is derived from patient-reported questionnaire and is used to measure lower back pain. Due to its accuracy, PROMIS PF can be a valuable metric to evaluate patient well-being following treatment and assist in evidence-based decision-making for treatment options for patients with spinal conditions.¹²

Domain Goals and Tasks

This project emerged from a collaboration between two computer scientists with four medical researchers from the Orthopedic Research Center and the Department of Population Health Sciences at the University of Utah. The domain scientists are currently investigating the use of PROMIS scores as a measure of patient well-being and progression of PF following various procedures for spinal ailments. In this project, we specifically target treatment options for intervertebral disc herniation. In meetings on a biweekly basis over 18 months, we collected notes on current EHR and PROMIS score use within the Orthopedic Research Center to identify domain goals and inform the design of our tool.

Two of the collaborators are spinal surgeons who have not used visualization of EHR data when considering a patient's options for treatment. Instead, their assessments have been based on past experiences. When determining patient treatment options, they take into account demographics, medical comorbidities such as diabetes, prior treatments, and current symptoms and severity. They then choose the treatment that is likely to result in the best outcome while also considering other factors such as recovery time and cost. The main treatment options considered by our collaborators for patients suffering from intervertebral disc herniation are hemilaminectomy (a surgical procedure), steroid injection, and physical therapy, as well as their combinations thereof. Because the medical histories and collected EHR data for the patient population are extensive and involve a variety of records and data types, we sought to develop a visual analysis solution that combines our collaborators' data into a comprehensive dynamic interface that helps them identify trends in patient outcomes. We identified three functionality requirements that inform the design of Composer, defined below:

- R1. Define meaningful cohorts of patients and analyze how this subset of patients reacts to various treatments and procedures. The clinicians need to be able to form cohorts from the EHR database on patient demographic information, treatment history, medical records, and initial PF scores.
- R2. Compare the outcomes of different cohorts, for example, PF outcomes following different treatment options in otherwise identical cohorts, or to identify an effect of a comorbidity.

R3. Normalize PFscores in several ways to successfully analyze and compare cohort outcomes, following an event, such as surgery.

Related Work

Visualization of patterns in patient medical histories helps identify risk factors that influence patient recovery following treatment.¹³ Recently developed clinical tools provide visual support for users, often in the form of aggregated representations of patient data derived from EHR as well as visual comparisons for patient outcomes and trajectories.^{5,7,14} Composer is related to various tools and techniques for cohort definition and EHR analysis, which we discuss below.

Cohort Definition

Cohort definition is a vital first step for analysis. Emergent patterns identified in cohort behavior and outcome remain dependent on the accuracy of the cohort creation,¹⁵ and therefore, cohort definition tools often provide visual feedback to track stages in cohort definition.¹⁵ We included a visual representation of each filter layer for a cohort in Composer and have extended this idea to allow dynamic changes to filters.

Cohort Comparison

Current visual tools often provide users the ability to compare clinical pathways and outcomes of patients. These comparisons help users identify differences in patient outcomes between two defined cohorts and diverging event sequences within a given cohort's records.⁵ Normalization to a standard time metric and alignment at events in the patient histories facilitate comparison and highlight patterns within the data.¹³ This time metric, often in the form of days or visits, allows patient histories to be viewed along a common axis. A tool developed by Bernard et al¹⁴ allows realignment of events, e.g., when metastases develop in cancer patients. By sorting and realigning, users can better see trends between events and their corresponding phases. Comparisons can be used for identifying both significant differences as well as similarities and recurring patterns. In contrast to Bernard et al, Composer represents patient trajectories as single lines layered over one another, which allows visualization of a larger number of patient trajectories at once. In Composer, we normalize patient data to a standard day metric and allow users to realign scores to a common-procedure event. This facilitates comparison of score fluctuation for cohorts containing several hundred patients after given events by viewing the patient score change aligned on a common axis.

Aggregation

Much patient data include event sequences and temporal information. With a large amount of patient data over a span of years, visualization of patient care pathways and events can prove difficult. Clinicians must be able to identify patterns of events within a single patient's medical history and recurring trends between multiple patients' records.¹⁶ Data, therefore, are often aggregated and summarized to identify emergent patterns within the cohort's medical timelines and track progression.¹⁷ Aggregation can help

with pattern identification within complex temporal data by reducing the visual complexity, although it can also hide subtle trends in the data.^{16,18} Composer uses aggregation of individual scores to show emergent trends in PROMIS score fluctuation without the clutter of hundreds of individual plotted trajectories of patient scores at once. Users can view the scores individually or aggregated at their discretion.

Making Relationships in the Data Explicit

Many recent tools facilitate cohort definition and analysis by making relationships between events and static attributes more explicit. Bernard et al's visual analysis tool for patients with prostate cancer visualizes distributions of static attributes in the data and indicates when an attribute's frequency is higher or lower in the cohort relative to the population. This visual information is valuable to the domain expert as it provides insight into filter constraints on attributes that might have influenced a subset of patient outcomes.¹⁴ Du et al's EventAction is a prescriptive visual tool for event sequences. It provides plots showing positive and negative correlations between categories and outcomes.¹⁹ Another method of highlighting significant relationships within the cohort data is through visual hierarchy and color. Many visual tools provide color-coded highlighting to emphasize significant events.^{14,20} By making these relationships explicit, users can make informed decisions to determine the next steps. We have incorporated these methods in Composer by providing distribution plots to show the number of patients in the entire population who meet the requirements for each filter category. For example, users can see the distribution spread of patient body mass index measurements. We also provide visual representation of each filter constraint on a given cohort along with the number of patients at each filter stage.

Composer Design

Composer, shown in ~Fig. 1, consists of two components: the cohort definition interface and the visualization of PROMIS PF scores. The cohort definition interface is contained within the collapsible sidebars on the left, while the outcome score interface is placed on the right. We chose to encode the score trajectories as a line plot, similar to the style of chart our collaborators currently use to represent PROMIS score trajectories, as this is both perceptually efficient and a common representation to view change in a metric over a period of time.

Cohort Creation

Our collaborators need the ability to define a cohort from a set of specific attributes and medical histories (R1). In Composer's filter sidebar (see ~Fig. 1A–E), cohorts can be defined by demographic information such as age or gender, in addition to other factors deemed relevant, like smoking habits. The filter sidebar is divided into demographic, score, and CPT (current procedural terminology; codes used to identify procedures) sections. Within the demographic filters, we use histograms to visualize the distributions of attributes in the patient population (~Fig. 1C). The histograms also serve as means to interact with a filter through

Fig. 1 Composer overview. Composer consists of interfaces for flexibly defining cohorts, and for displaying the physical function scores of patients treated for back problems over time in these cohorts. (A) Patient cohorts can be added and branched in the cohort control interface. (B) A history of all filters applied to the selected cohort. Cohorts can be defined using (C) filters applied to demographic information, (D) recorded score frequencies, (E) and presence or absence of procedural codes. (F) The main interface is a chart showing either individual lines or aggregated areas. A zero-point for the PROMIS scores, indicated by the horizontal red line, can be flexibly defined to align all patients by a specific event, such as a medical intervention. (G) The layer panel provides the ability to hide layers corresponding to the cohorts. (H) Users can select individual patient lines to show orders associated with their medical records in the timeframe specified in the timeline below the main plot. Selected patients are identified by their patient id, shown on the left-hand side of the event line.

brushing for quantitative attributes and selections for categorical ones. In addition to demographic variables, cohorts can also be defined by the number of recorded PROMIS scores for a patient (►Fig. 1D), or based on the presence or absence of procedure codes in patient histories (►Fig. 1E). This allows analysts to, for example, separate patients that have received a specific surgery from those who have not. With each cohort refinement, a filter layer is added to the sidebar as a visual history of filters used and cohort size at the given filter (►Fig. 1B). Individual filters and cohorts can be removed from the filter history or updated at any time in the cohort sidebar (►Fig. 1A). Composer enables analysts to define multiple cohorts simultaneously. Each cohort is represented as a colored bar and assigned a unique label and color, which is kept consistent across the interface. Within the bar, filters are represented as white nodes. If more than three filters are present, they are aggregated.

To facilitate cohort comparison (R2), cohorts can be branched. Once branched, the filter constraints of the parent cohort are duplicated in the branch but can be refined independently. This allows users to add diverging filters for an attribute that an analyst believes may influence the outcome of a treatment. For example, users may want to see if there is a difference in patient trajectories after physical therapy, if they have also had a steroid injection. To do that, they can define an

initial cohort, branch it, and apply filters for subsequent steroid injections versus no injections to the branches.

Outcome Score Comparison

PROMIS PF scores for the defined cohort are visualized as individual lines showing the course of PF for each patient over time. The timewindow can be resized as desired. By default, we align by the first PROMIS score, yet alignment by a specific clinical event, such as surgery or the start of physical therapy, is often more informative. When different cohorts are aligned by different events this way, the relative progression after the event can be evaluated. This facilitates comparison between cohorts (R2) by allowing the user to manipulate the alignment and scale in a dynamic way (R3). We use juxtaposition and superimposition to compare between cohorts,¹⁸ which have different trade-offs as far as required display space and clutter in a single plot are concerned. Juxtaposition allows users to add multiple plots to evaluate cohort trajectories in a side-by-side comparison (►Fig. 2). Superimposition shows different layers on top of each other (►Fig. 1F). We allow analysts to toggle layers individually (►Fig. 1G).

Dynamic Score Scales and Normalization

The PF scores used by the domain experts are often subtle in absolute measured change (see ►Fig. 3A), yet these

Fig. 2 Differences in PROMIS scores after surgery and injection compared by (A) layering and (B) juxtaposition of multiple plots. Both methods allow for comparison of score change after different treatment events. (A) Treatment options in layers. (B) Juxtaposition in multiple plots.

Fig. 3 View of score plots using (A) absolute and (B) relative scales. Each line represents an individual patient. Relative scales show change in PROMIS PF score, calculated from the score at the day zero event. In this case the patient score trajectories are aligned by the day of surgery. With a larger cohort, the general trend for patient progression can be difficult to see, which we address by providing aggregation functionality. (A) Absolute score scale. (B) Relative score scale.

subtle changes often have significant impact on the perceived well-being of patients. Changes in patient scores are further obscured as patients in the same cohort have different baseline scores. To emphasize change and normalize the baseline, analysts can view scores on a normalized scale that visualizes relative score change for the patients, as shown in ►Fig. 2B. With the option of both absolute and relative score scales, analysts can assess the cohort's overall trend in baseline score measurements as well as trends in score fluctuation. By showing relative score change and making the relationship between cohort scores more explicit, analysts can see differences in outcome trajectories during comparison more clearly. In addition, users have the ability to adjust the timeframe of the line chart. The timeframe is specified through brushing a selection of the lower timeline that extends the minimum and maximum range of days for all patient records (see ►Fig. 3).

Separation of Scores by Quantiles

Even in a well-defined cohort, patient outcomes can be markedly different. Due to this heterogeneity, our collaborators need the ability to separate the cohort into quantiles that communicate how, for example, the PF changes for the top 25% of patients in the cohort (see ►Fig. 4A). In Composer, a cohort can be divided by quantiles. We calculate these quantiles by the average change in score over a user-adjustable period of days following a given event.

Aggregation of Scores

Frequently, our collaborators do not need to view individual patients, but rather are interested in aggregate representation of scores. To address this need, we provide means to aggregate the scores of a cohort to visualize the interquartile range with a line representing the median. Aggregated cohort scores can also be separated by quantiles to more clearly identify any difference in score change within subsets of the cohort that have different baseline measurements, as shown in ►Fig. 4B.

Individual Patient CPT History View

For further analysis of procedure code distributions and procedure frequency, analysts can select an individual patient from a group of patient trajectories in the score chart to view all orders associated with that patient's medical history (see ►Fig. 1H). These histories are cropped to the timeframe specified in the score chart and aligned with its timeline. For example, if the score chart shows trajectories between 20 days before an injection and 60 days after, the individual timeline would reflect the same timeframe. These events can provide context for individual cases, but can also be used to further filter a cohort. Analysts can view patient histories by selecting the patient's PROMIS scores on a given plot. The events then appear below the plot, aligned on the same time.

Implementation

Composer is open source and was developed with TypeScript using the D3.js library for visualization. The prototype is a Phovea client/server application.²¹ The code for Composer can be found at <https://github.com/visdesignlab/Composer>. Data used for development and to inform the usage scenario were sourced from a sample of EHR provided by our collaborators from the Orthopedic Research Center's database and were preprocessed in Python.

Usage Scenario

Here we describe a usage scenario to illustrate a typical use case for composer as it can be used by our domain collaborators.

A surgeon sees a patient suffering from a herniated disc. While evaluating potential treatment options for the patient, she defines a cohort in Composer using constraints based on the given patient's medical history. She filters by the patient's age range, specifies the cohort to only include diabetic patients, and filters just those patients that have had physical therapy evaluation. The cohort defined by these patient-specific filters contains 3,317 patients. She branches the cohort and filters the initial branch by those that have

Fig. 4 View of patient scores separated and color coded by quantiles. The PROMIS PF scores were separated into quartiles, shown as individual lines in (A) and aggregated area charts in (B). The orange marks represent the top quartile, the yellow marks the interquartile range, and the blue marks the bottom quartile. (A) Quantiles color coded. (B) Quantiles aggregated.

had surgery, but have not had an injection. She then filters the secondary branch by those patients that have had an injection but not surgery. Aligning each cohort by the surgery or injection event they were filtered by, she can view the diverging cohorts superimposed over one another and visually compare differences in PROMIS PF score fluctuation between the two. She can then aggregate the individual scores to show only the median PROMIS score within the cohort. Next, she normalizes the PROMIS scores from the absolute score measurement to relative score change, so that she can visually compare the difference in score change between the two to determine what treatment appears to produce better outcomes (→Fig. 2). After comparing the change in score across a span of 150 days after treatment, she can see that surgery had a greater positive change in PF, which is clearly visible after the first month (→Fig. 2A). She can take this into consideration when determining patient treatment options, and show this visualization to the patient when discussing treatment options.

Discussion and Limitations

Composer is under active development, with progressive iterations being made in response to feedback received from meetings with collaborators.

Evaluation: We considered various strategies to evaluate our contribution, including collecting feedback from our collaborators, and comparing to other tools. While we have received positive feedback from our collaborators, we chose to not report it in detail due to the potential for biases. Ultimately, we have chosen to validate Composer through a usage scenario and the careful justification of our design decisions, which are accepted practices in user-centered design.²² However, the larger question is whether using a tool like composer will lead to better outcomes. We are currently planning a longitudinal study using the tool and measure provider and patient satisfaction, but also outcomes. However, such a study is beyond the scope of this article.

Data integration: Currently, the data used in Composer are a large but static dataset of patients pulled from the Orthopedic Center's database. By using a static snapshot, we have full control over processing and data manipulation for initial development while avoiding issues such as permissions and compatibility associated with a deep integration with the EHR system. We expect to be able to run a longitudinal evaluation without integrating Composer; however, this creates manual effort when incorporating new patient data or updating existing data. As we develop Composer beyond its proof-of-concept stage and past a formal evaluation, we intend to integrate the tool with our collaborator's EHR system.

Data cleanup: A challenge common to systems operating on data extracted from EHRs is the data's messiness and inconsistency. We address sparse outcome scores by interpolation, yet we acknowledge the limitation in accuracy for interpolated patient trajectories for those patients that have lower score frequencies. We exclude patients with fewer

than three PROMIS PF score. We also do not currently consider systematic biases in score trajectory: for example, it is likely that we have less data for patients with good outcomes, as they do not come for follow-ups. We hope to mitigate these limitations in future iterations of the tool by making uncertainty in patient trajectories more explicit in visualization and statistical representation.

Conclusion and Implications for Future Work

In this article, we outlined the domain analysis and the design of Composer, an application to visualize and compare patient cohorts and their PF trajectories. This tool was developed in collaboration with domain experts from the Orthopedic Research Center at the University of Utah, with their current research in the efficacy of PROMIS scores to evaluate PF of patients with lower back conditions. Immediate development of the tool will focus on addressing the limitations described in the previous section. In the near future, we plan to provide a more extensive statistical breakdown of cohort medical history with the inclusion of International Classification of Diseases (ICD) codes. As distributions of events and attributes become more explicit, users will be able to apply more accurate filtering constraints to define cohorts. Additionally, we plan to provide more control of the CPT filter codes as they appear within the patient record, and inclusion of sequence-specific event filters. As recent literature has shown that medical event sequences can provide important clues on patient outcomes,^{5,8,19} Currently, target patient outcomes are interpreted implicitly by evaluating score trajectories of a body of similar patients. We intend to improve interpretation of target patient outcomes through explicit data-driven forecasting of score trajectories using a larger patient sample, informed by previous work from Buono et al.²³ Composer's initial development targets orthopedic patient comparisons and evaluation, and we expect to be able to generalize it to other cases where outcome measures over time are the subject of the analysis. We also anticipate that our cohort definition interface could be applied in an even broader context.

The long-term goal for Composer is the addition of an interface for shared decision making in which insight from exploration in the current interface could be translated into visualizations that would facilitate the explanation of treatment choices and potential outcomes to the patient, and the integration of other measures, such as cost. As previously mentioned, we also plan a clinical evaluation of the tool.

Clinical Relevance Statement

Visual cohort analysis has gained attention as an informative analytic tool in healthcare with its potential to help clinicians assess optimal treatment options for patients with preexisting conditions that can influence recovery and treatment.

Multiple Choice Questions

1. What is the benefit of patient score normalization in visual cohort analysis?
 - a. Normalization to a standard time metric and alignment at events in the patient histories facilitate comparison and highlight patterns within the data.
 - b. By normalizing to show relative score change, we can make the relationship between cohort scores more explicit and differences in outcome trajectories during comparison clearer.
 - c. Both a and b.
 - d. None of the above.

Correct Answer: The correct answer is option c.

2. What are the primary treatment options for the sample patients used in this work?
 - a. Surgery.
 - b. Steroid injection.
 - c. Physical therapy.
 - d. All of the above.

Correct Answer: The correct answer is option d.

Protection of Human and Animal Subjects

Our work does not involve any studies with human or animals performed by any of the authors.

Funding

This project was funded by the University of Utah Orthopedic Research Center and NSF IIS 1751238.

Conflict of Interest

None declared.

References

- 1 Lee J, Maslove DM, Dubin JA. Personalized mortality prediction driven by electronic medical data and a patient similarity metric. *PLoS One* 2015;10(05):e0127428
- 2 Gallego B, Walter SR, Day RO, et al. Bringing cohort studies to the bedside: framework for a 'green button' to support clinical decision-making. *J Comp Eff Res* 2015;4(03):191–197
- 3 Cella D, Riley W, Stone A, et al. The Patient-Reported Outcomes Measurement Information System (PROMIS) developed and tested its first wave of adult self-reported health outcome item banks: 2005–2008. *J Clin Epidemiol* 2010;63(11):1179–1194
- 4 Thadhani R, Tonelli M. Cohort studies: marching forward. *Clin J Am Soc Nephrol* 2006;1(05):1117–1123
- 5 Perer A, Wang F, Hu J. Mining and exploring care pathways from electronic medical records with visual analytics. *J Biomed Inform* 2015;56:369–378
- 6 Du F, Plaisant C, Spring N, Shneiderman B. Finding similar people to guide life choices: challenge, design, and evaluation. Paper presented at: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems; Denver, Colorado; May 06–11, 2017. New York, NY: ACM; 2017:5498–5544
- 7 Gotz D, Wang F, Perer A. A methodology for interactive mining and visual analysis of clinical event patterns using electronic health record data. *J Biomed Inform* 2014;48:148–159
- 8 Franklin L, Plaisant C, Rahman KM, Shneiderman B. Treatment explorer: an interactive decision aid for medical risk communication and treatment exploration. *Interact Comput* 2014;28(03):238–252
- 9 Gruber-Baldini AL, Vellozo C, Romero S, Shulman LM. Validation of the PROMIS® measures of self-efficacy for managing chronic conditions. *Qual Life Res* 2017;26(07):1915–1924
- 10 Houck J, Wise Z, Tamanaha A, et al. What does a PROMIS t-score mean for physical function? *Foot Ankle Orthop* 2017;2(3). Doi:2473011417S000200
- 11 Hung M, Hon SD, Franklin JD, et al. Psychometric properties of the PROMIS physical function item bank in patients with spinal disorders. *Spine* 2014;39(02):158–163
- 12 Brodke DS, Goz V, Voss MW, Lawrence BD, Spiker WR, Hung M. PROMIS PF CAT outperforms the ODI and SF-36 physical function domain in spine patients. *Spine* 2017;42(12):921–929
- 13 Wang TD, Plaisant C, Quinn AJ, Stanchak R, Murphy S, Shneiderman B. Aligning temporal data by sentinel events: discovering patterns in electronic health records. Paper presented at: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems; Florence, Italy. New York, NY: ACM, 2008:457–466
- 14 Bernard J, Sessler D, May T, Schlomm T, Pehrke D, Kohlhammer J. A visual-interactive system for prostate cancer cohort analysis. *IEEE Comput Graph Appl* 2015;35(03):44–55
- 15 Krause J, Perer A, Stavropoulos H. Supporting iterative cohort construction with visual temporal queries. *IEEE Trans Vis Comput Graph* 2016;22(01):91–100
- 16 Monroe M, Lan R, Lee H, Plaisant C, Shneiderman B. Temporal event sequence simplification. *IEEE Trans Vis Comput Graph* 2013;19(12):2227–2236
- 17 Widanagamaachchi W, Livnat Y, Bremer P-T, Duvall S, Pascucci V. Interactive visualization and exploration of patient progression in a hospital setting. *AMIA Annu Symp Proc* 2017;2017:1773–1782
- 18 Gleicher M. Considerations for visualizing comparison. *IEEE Trans Vis Comput Graph* 2018;24(01):413–423
- 19 Du F, Plaisant C, Spring N, Shneiderman B. Eventaction: visual analytics for temporal event sequence recommendation. Paper presented at: 2016 IEEE Conference on Visual Analytics Science and Technology (VAST). IEEE; 2016:61–70
- 20 Gschwandtner T, Aigner W, Kaiser K, Miksch S, Seyfang A. Design and evaluation of an interactive visualization of therapy plans and patient data. Paper presented at: Proceedings of the 25th BCS Conference on Human-Computer Interaction. British Computer Society; 2011:421–428
- 21 Gratzl S, Gehlenborg N, Lex A, Strobel H, Partl C, Streit M. Caleydo Web: an integrated visual analysis platform for biomedical data. Paper presented at: Poster Compendium of the IEEE Conference on Information Visualization (InfoVis '15). IEEE; 2015
- 22 Greenberg S, Buxton B. Usability evaluation considered harmful (some of the time). Paper presented at: Proceedings of the SIGCHI Conference on Human factors in computing systems. ACM; 2008: 111–120
- 23 Buono P, Plaisant C, Simeone A, et al. Similarity-based forecasting with simultaneous previews: a river plot interface for time series forecasting. Paper presented at: 2007 11th International Conference Information Visualization (IV'07). IEEE; 2007:191–196