Monitoring and Improving Quality of Care with Interactive Exploration of Temporal Patterns in Electronic Health Records

Phuong D Ho, MD, Taowei David Wang, PhD[1, 2], Krist Wongsuphasawat, Catherine Plaisant, PhD[1] Ben Shneiderman, PhD[1, 2], Mark S. Smith, MD[3], David Roseman, MHA[3]

[1]Human-Computer Interaction Lab, University of Maryland, College Park, MD;
[2]Department of Computer Science, University of Maryland, College Park, MD;
[3]ER One Institute, Washington Hospital Center, Medstar Health, Washington, DC;

Abstract *(150 words unstructured)*

As the use of electronic health records (EHRs) spreads, interactive query interfaces will assist researchers and clinicians to quickly find historical records that include specific temporal patterns. This paper describes such an interface called LifeLines2, and summarizes how it was applied to a set of clinical problems. Our experience indicates that LifeLines2 was helpful to: 1) quickly identify unknown data quality issues, 2) speed up the finding of patients exhibiting specific temporal patterns of interest, 3) allow repeated queries to facilitate trend analysis and comparisons that would not have been conducted otherwise for lack of human resources. Our case studies illustrate the potential impact on patient safety, quality assurance and process improvement.

1. INTRODUCTION

As the use of electronic health records (EHRs) spreads, there are growing opportunities for their use in patient care, clinical research, quality assurance, and alarm specification. For example, setting an alarm for patients on heparin with a precipitous drop in platelet counts requires specificity around the definition of "precipitous." By querying existing EHR databases, physicians designing the alarm can iteratively test the logic of the alarm (*e.g. absolute drop* or *relative drop*? in what *time range*?) and validate it with a large amount of data. In another example, researchers might start by looking for "all patients who were discharged from the emergency room then admitted again within a week", or study the incidence of patient "bounce-back" by looking for specific temporal sequences of room transfers.

One commonality in these sophisticated medical queries is that they often have a temporal component. However, specifying temporal queries in SQL is difficult even for computer professionals specializing in such queries. Database researchers have made progress in representing temporal abstractions and executing complex temporal queries [1, 2, 3, 4], but there is very little research that focuses on making it easy for clinicians and medical researchers to both specify the queries and examine results visually. Currently available graphical user interfaces make it possible to specify the simpler conjunctive queries such as "find patients who transferred to ICU rooms and the general floor rooms (in any order)" [5, 6]. However, this leaves users with the burden of examining large numbers of results in search of matching temporal sequences (*e.g.* transferred to ICU and then *subsequently* to floor).

We believe interactive query interfaces that allow physicians to explore data that have specific temporal patterns will dramatically increase the benefits of EHR databases. In addition, comprehensive presentation of the search results can help users see patterns, outliers and data quality problems. This paper describes the interface developed by the University of Maryland (UMD) called Lifelines2 (Figure 1) and summarizes how it has been applied to a set of clinical problems investigated by researchers at the Washington Hospital Center (WHC). Our experiences indicate that Lifelines2 is helpful to: 1) quickly identify unexpected data quality issues, 2) speed up the finding of patients exhibiting specific temporal patterns of interest, 3) allow repeated queries to facilitate trend analysis and comparisons that would not have been conducted otherwise due to the lack of human resources. Our case studies illustrate Lifelines2's potential impact on patient safety, quality assurance and process improvement.

2. BACKGROUND

Handling time-related concepts is essential in medicine. A survey [3] lists the many applications that use the history of temporal events, and review the latest contributions to time-aware decision support systems. While causality, natural language, and argumentation still hot topics in the AI and Medicine research community, Augusto observes that there are still opportunities for providing much needed search tools. Some research systems provide temporal access languages to support limited visual queries from end-users [7, 8, 9, 10, 11, 12], but many of these suffer the same accessibility difficulties of SQL or they require an understanding of the underlying database structure.

The large body of related computer science work can be grouped into three general areas: time theory, databases, and visualizations. None of the systems discussed in those papers enable query or visualizations of patterns across multiple histories (*e.g.* temporal cross-patient queries) that are essential in EHR systems. This paper presents the case studies of Lifelines2, which combines the power of temporal queries with graphical visualizations, and reports its benefits and limitations.

Time Theory: Much of the seminal work in computer science relating to time [13, 14, 15] stems from artificial intelligence, time reasoning, and early natural language processing.

Databases: Due to the complexity of formulating SQL queries, several approaches have made database query more accessible to a broader spectrum of users, *e.g.* Query By Example (QBE) [16], the visual query mechanism used in Microsoft's Access, TSQL[10], a hybrid between QBE and Extended Entity-Relationship diagrams [17, 18]. MQuery [19] targets various types of streaming data. PROTEMPA allows the interactive specification of temporal abstraction [20].

Visualizations: Chittaro and Combi proposed three alternative visual metaphors for querying temporal intervals [21]. Hibino and Rudensteiner introduced a direct manipulation Temporal Visual Query Language (TVQL)[12] to support Allen's thirteen relational primitives [14]. Interestingly, none of the work described above address the visualization of the returned results. However, applications such as TimeSearcher [22], DataJewel [23], KNAVE [9] and LifeLines [24] offer visualizations that cluster results and highlight temporal patterns. LifeLines provides a compact hierarchical timeline visualization for personal histories organized by facets, such as doctor visits, lab tests, and medications. It focuses on a single record and does not offer a query mechanism for discovery across multiple records. Many systems have built on LifeLines, e.g. [25]. Shahar proposes a Knowledge Based Temporal Abstraction model RÉSUMÉ [1] for a single patient.

Recently there has been more interest in developing tools that enable queries and visualizations of patterns across multiple patient records. The implementation of PatternFinder in Amalga [26] builds on an early prototype PatternFinder [27]. In these systems, a limited set of temporal filters could be chained together with test value changes relative to the previous event. Klimov *et al.*, improves upon Shahar's designs and created a system that performs temporal abstraction and visualization for multiple patients [28, 29]. Instead of using only exact matches in search, Similan explores similarity search in temporal medical domains [30]. While these systems contribute to different scenarios and tasks physicians and researchers may find useful, they have largely been evaluated in a limited situations such as in controlled experimental settings, or through only a single case study. LifeLines2, on the other hand, has been applied to eight different medical scenarios [31, 32, 33], and these case studies offer a deeper look into how physicians would utilize such an interactive visualization tool.

3. DESIGN OBJECTIVES

Developers from UMD and physicians from WHC worked together on three initial medical case studies and identified four common characteristics in temporal queries. First, there is often an important event (*sentinel event*) with respect to which all other events are examined (*e.g.* exposure to heparin in heparin-induced-thrombocytopenia). Secondly, event sequences are important (*e.g.* baseline reading of normal platelet count followed by administration of heparin followed by low platelet count). Thirdly, parts of the sequence may require temporal range (*e.g.* low platelet reading occurs within 5-9 days after the heparin event). Fourthly, event prevalence over time for a set of patients is important. These search characteristics are currently not supported in current EHR systems.

Over a two-year long iterative design process, we designed and implemented a timeline-based visualization of patient records to help physicians review patient records. We also designed a number of interaction techniques. *Align*, *rank*, and *filter* operators in Lifelines2 [31] emphasize temporal ordering of events. T*emporal summaries* [32] accentuate event prevalence over time. Temporal summaries dynamically aggregate multiple events in user-specified granularities (year, month, day, hour, minutes, etc.) for the purpose of spotting trends over time and comparing several groups of records.

Lifelines2 focuses on only point categorical data. No numerical or interval data are handled. Every data point has a time stamp and an event name (category). These limitations can be overcome by first preprocessing the data (*e.g.* bracketing numerical data into high, normal, low categories).

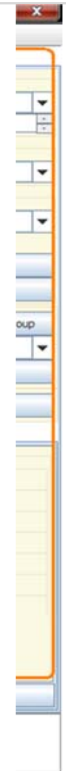

4. SYSTEM DESCRIPTION

Lifelines2 enables users to visually explore a collection of patient records (Figure 1(a)). Each record is represented as a horizontal strip, identified by its ID on the left, and stacked vertically. Events appear as triangles on the timeline, colored by their category. Mousing over an event reveals details of that event in a tooltip.

The control panel on the right (Figure 1(c)) allows users to manipulate the records. Users can select an event type by which to *align* all records. For example, Figure 1 shows all records aligned by their 1st ICU event. When alignment is performed, the timeline becomes relative to the alignment event.

Users can *rank* and *filter* records by the number of occurrences of an event type. Users can additionally filter by a temporal sequence. Finally, users can select records that contribute events to the *temporal summary* by drawing selection rectangles (Figure 1(b)). Records that fit these filter criteria are highlighted in orange gradient. Users can then decide to *keep* or *remove* these selected records.

For example, to find patients who were admitted, then transferred to the ICU on the first day and transferred to IMC (intermediate medical care) on the fourth day, users would first *align* all the records by ``Admit", display the *temporal summary* of ``ICU", then select the bar that represents the all "ICU" events on the 1^{st} day in the *temporal summary* and click on ``Keep Selected." Then users would show the distribution of ``IMC" and draw a box to select the fourth day and ``Keep Selected."

A similar process can be used for consecutive interval specification using different alignments and filtering. Lifelines2 is a Java application, utilizing the Piccolo 2D graphics framework [34].

5. STATUS REPORT

We identified a set of medical questions WHC researchers are interested in and conducted case studies based on them. We present two in detail here. The first case study looks at patient transferring out of a high-level care room to a lower level one, but only to bounce back to a higher level room soon after. The second case study looks at pulmonary patients whose conditions worsen after receiving BiPAP. The remaining four case studies are only succinctly described but demonstrate the diversity of questions that Lifelines2 can address.

All case studies are conducted following the same protocol. UMD developers obtain and preprocess de-identified data obtained from WHC and load them into Lifelines2. Developers and WHC physicians then collaboratively meet to review and explore the data. As new questions arise, new queries are executed within Lifelines2 interactively. Software bugs were sometimes identified during the process and fixed. Each case study typically takes 4-8 meetings, 1-2 hours per meeting over 2-3 months.

A. Case Study -- Bounce Back

Physicians look for patients who are discharged from a higher-level-care room to a lower-level-care room, and then back to a higher-level-care room within a small time frame. For example, ICU->Floor->ICU will be one such sequence. A temporal constraint is required to limit the query so that the second ICU is within 48 hours of the discharge from the first ICU (coinciding with the entering of Floor) to ensure that the bounce-back is likely related to the first ICU stay. This kind of event sequence may be indicative of a patient being transferred from a higher-level-care room prematurely. The set of sequences that describe various bounce-back patterns is listed below. IMC denotes the intermediate medical care and holds patients whose conditions are not severe enough to be in ICU but more severe than those Floor holds. MS is an overflow area for patients who are transferred to ICU when ICU is full.

ICU->Floor->{ICU or MS}

ICU->Floor->IMC

IMC->Floor->IMC

IMC->Floor->{ICU or MS}

ICU->IMC->{ICU or MS}

Current Method

Prior to 2007, a database administrator would query the database to pull out patient location data as an Excel file. This file would then be sent to the ICU director who would manually review the spreadsheet row by row and look for bounce-back patient. This process would take approximately 40 hours to process 3 months worth of data. Since 2007, the ICU director asked Ho (the primary author) to develop an automatic way to analyze the spreadsheet. After spending approximately 30 hours, Ho was able to develop a set of formulas in Excel to perform this task (Figure 2). Using these formulas, it would take about 10 minutes to process the same amount of data. However, there were known inaccuracies with using the formulas as it is difficult to account for unusual temporal ordering of events in Excel. Ho also attempted to perform the analysis in SQL but quickly came to the realization that it would be challenging to generate the appropriate queries [35] even though his undergraduate degree was in computer science. In the end, the director decided to accept the small inaccuracies in Excel and its poor presentation of temporal relationships as a trade-off for timely analysis.

Using Lifelines2

Ho would first apply a sequence filter for each of the sequences above. He would then align the patient records by the middle event of the sequence and subsequently select those that have third event occurring within 48 hours of the alignment, using temporal summaries. The process was repeated for each of the sequences, and the patient result sets were merged at the end.

At each step, Ho was able to quickly scan the result to ensure accuracy and spot anomalies. For example, the developers and Ho noticed that some patients are sent to rooms designated as MS (overflow rooms for ICU) before entering ICU. These patients effectively experienced a bounce-back, but were not captured because of the MS designation instead of ICU. We later replaced ICU with {MS or ICU} in our sequence. This error was not handled in the Excel method. The spreadsheet display made it extremely difficult to detect this kind of error. It was not until we performed the case study in Lifelines2 that this error was discovered.

On the other hand, because we had to convert data into the Lifelines2 format, a few initial mistakes were made. For example, some specialized ICUs were not included (*e.g.* cardiovascular operation recovery room). This type of error is discovered only by careful review of all the rooms and the conversion script. We had since fixed these problems.

We were able to find that for July-September 2009, there were 29 patients out of 576 (5.0%) that exhibited this bounce-back pattern. These rates are within acceptable ranges. The analysis took less than 30 minutes to complete.

Comparison

Identifying bounce-back patients with Lifelines2 was an order of magnitude faster than manual review. Taking out the initial time to develop the formulas, lifelines2 takes longer than Excel. However, there are three main advantages using Lifelines2. First, the features in Lifelines2 are able to capture the bounce-up pattern while the Excel method contained errors. Secondly, Lifelines2 allows physicians to visually verify the exactness of a sequence filter (*e.g.* a filter is too broad or too narrow). The visual representation facilitates the detection of strange patterns (*e.g.* transfer to ER after admission to ICU) or erroneous data entries that must be remedied before analysis.

Finally, physicians are able to modify an existing sequence filter quickly by themselves in Lifelines2. The sequence filter feature, alignment, and temporal summary serve as simple and effective tools for physicians to specify the temporal sequences they are interested in. Selecting the right sequence, temporal range, and alignment is easier than specifying formulae in Excel or constructing SQL queries. Transferring Lifelines2's utility to other medical scenarios can also be much faster.

Next Steps

The Lifelines2 bounce-back study is repeated for October-December 2009 (5.4%) and January-March 2010 (3.7%). We are conducting this study on a regular basis to ascertain the quality of the results. If the quality continues to be good in comparison to manual review, Lifelines2 may potentially be adopted as the primary way to perform bounce-back studies to monitor quality within WHC.

B. Case study – BIPAP Usage

There are many reasons for a patient to have difficulty in breathing. When this problem becomes severe, physicians often have to intervene with mechanical support. In the most severe cases, this requires the insertion of a breathing tube (intubation) down the trachea followed by the tube being hooked up to a breathing machine or ventilator. This type of intervention is uncomfortable and can also cause harm if not monitored appropriately. In many cases, if the difficulty in breathing is serious but not severe, physicians will try to treat the underlying problem while providing less invasive breathing support via a machine called Bi-level Positive Airway Pressure (BiPAP). BiPAP is a machine that delivers forced oxygen into the patient via a facemask. While also uncomfortable, it is not invasive, and in some medical conditions BiPAP has been shown to be more effective in reducing mortality than intubation.

Prior to 2008, BiPAP was prescribed purely by physicians without the use of any literature-based protocols for use and manipulation. In late November 2008, WHC switched BiPAP usage to be protocol-driven. The physicians are interested in whether the new protocol reduces the rate of escalation of care or failure of BiPAP (*i.e.* requiring intubation after BiPAP, or going to a higher-level of care room) (Figure 3).

After collecting and merging data from both the pulmonary database and the patient room, there are a total 6583 pulmonary patients during the period of October 2008-February 2009. 643 patients had BiPAP in that period. We separated the records into pre-protocol (October-November 2008) and post-protocol (December-February 2009) groups. There are 280 BiPAP patients in the former and 463 in the latter. The temporal patterns indicative of escalation are listed below.

BiPAP->Intubation

BiPAP->ICU

No {Floor, IMC, or ICU}->BiPAP->IMC

No {IMC or ICU}->BiPAP-> IMC

Floor->BiPAP->IMC

Floor->BiPAP->{ICU or MS}

IMC->BiPAP->{ICU or MS}

Current method

WHC does not currently have a meaningful way to assess this. This would require a manual chart review of all 643 patients on BiPAP. Each chart review would take at minimum 3 minutes, leading to an exceedingly long evaluation time. It is possible that someone might take a subset for manual review and make generalized conclusions, but no one will continue it longitudinally as it is too time-consuming.

With Lifelines2

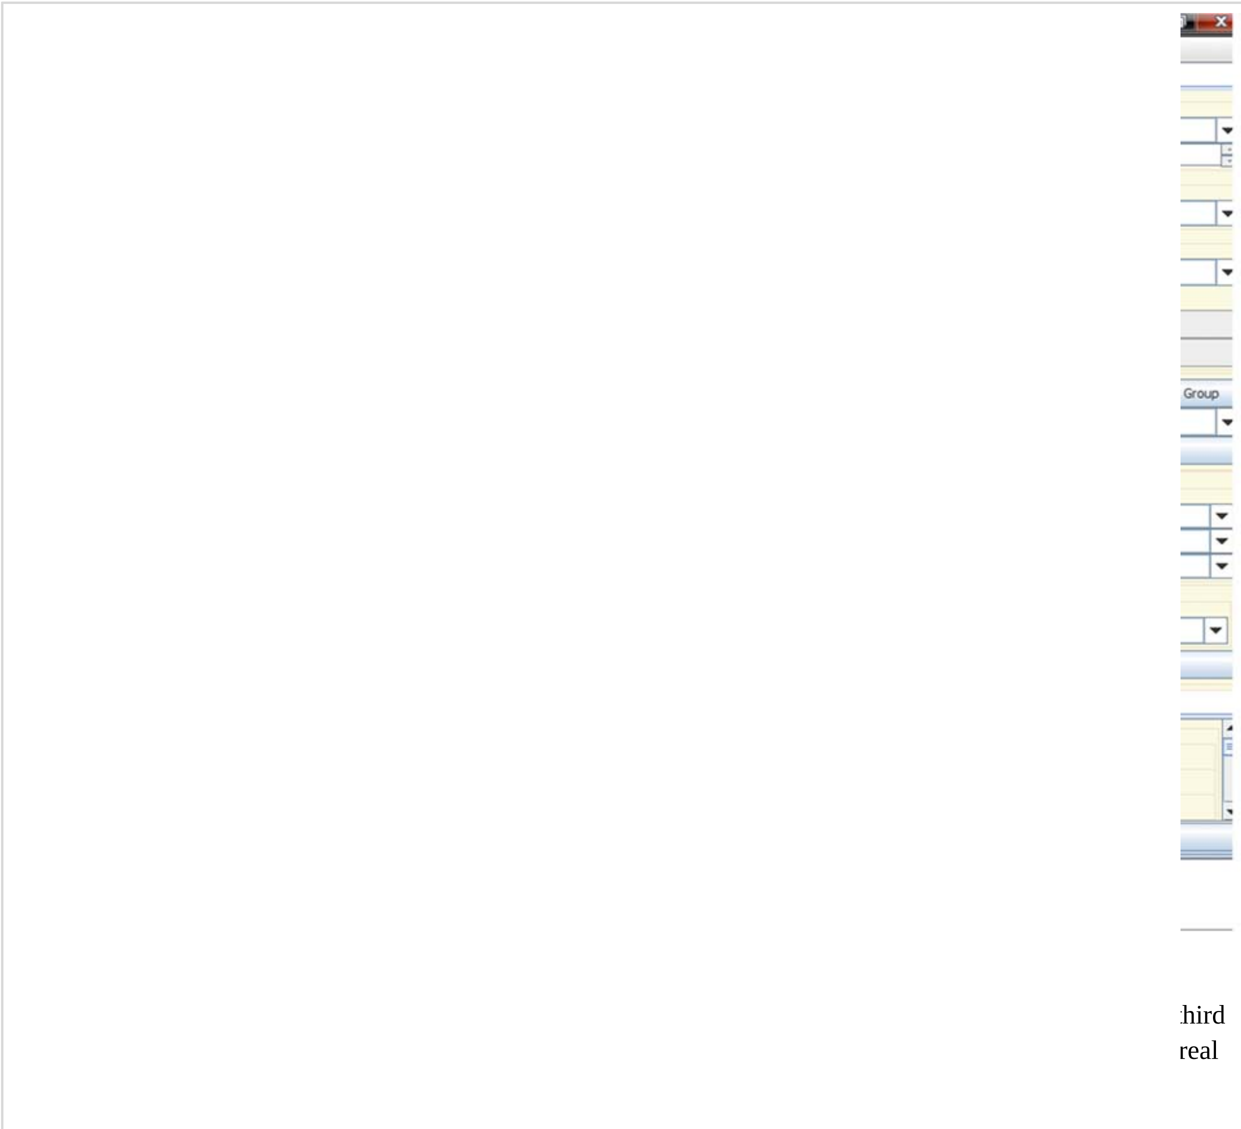

For each of the pre-protocol and post-protocol groups, we apply the sequence filters listed above. The results are then combined to represent a list of patients who potentially experienced escalation due to BiPAP. Exploration of the list reveals inconsistencies in the data. For example, physicians know that intubation is typically applied for at least 2 to 5 hours and only in ICU. When a patient record shows isolated intubation events or if intubation occurs outside of ICU, it is indicative of a data entry error. Secondly, extubation events indicate when a patient stops using intubation, however, extubation events are recorded inconsistently. We can infer random extubation events as an error (Figure 3). We can also infer the existence of extubation events when some other assisted breathing device is used. Because of these inconsistencies, physicians are required to review the results in Lifelines2 to determine whether escalation occurred.

After filtering, there are 99 potential escalation cases in the pre-protocol group and 128 in the post-protocol group. These potential cases are then reviewed by physicians in Lifelines2 to determine whether they are probable. In the end, we observe a large decrease in the escalation rate from 20% (56 out of 280) to 12.9% (60 out of 463).

In contrast to manual chart review, Ho was able to determine whether a patient had an escalation in 15 seconds on average.

Comparison
A traditional chart review would take physicians over 1,800 minutes to examine 600 BiPAP patients over five months. In contrast, Lifelines2 provides a fast and flexible search for relevant patients for review, and reduced the number of patients that require review from 642 to 228. Furthermore, Lifelines2's visual display expedites review by up to 12 times.

On the other hand, the Lifelines2 review process is not as detailed or nuanced as a chart review since we only use only pulmonary data and room data of patients while hospital charts contains the patients' complete history. However, Lifelines2 offers a reasonable alternative to chart review, and allows for ongoing surveillance of the effectiveness of the new protocol.

Next Steps
The reduction in escalation rate is promising despite its preliminary nature. WHC is currently preparing to perform chart reviews for the same BiPAP patients to get an accurate count of escalation cases. We intend to compare chart review result to Lifelines2 results to ascertain the fidelity of the Lifelines2 method. We fully expect discrepancies and flaws in the Lifelines2 analysis process. However, we are also optimistic that these issues can be corrected to ensure a correct and speedy analysis process.

C. Other Case Studies

In addition to the above two case studies, Lifelines2 has been applied successfully to search for patients who exhibit contrast-related nephropathy and heparin-induced thrombocytopenia [32]. We were able to use Lifeliens2 to replicate study linking daylight savings time shift to higher incidents of heart attacks [36] using WHC data in under 30 minutes. We also conducted a case study find patients who *Steps-Up* to higher level care room immediately after being admitted to a lower level room, indicating potential insufficient triage.

6. DISCUSSIONS

A. System Limitations

Lifelines2 is better suited for users who have a good idea about their query. Users who do not have a particular query in mind does not benefit as much. To further improve exploratory search for temporal data, we have experimented with two more prototypes. Similan uses similarity search to find similar event sequences instead of exact sequences [30]. LifeFlow automatically aggregate event sequences so frequent sequences and rare outliers can be spotted easily. We are integrating all of these features in one single tool.

Lifelines2 only handles temporal categorical data. While we can deal with numerical data types (*e.g.* blood pressure reading) or attribute data (*e.g.* gender, age) by preprocessing, this significantly limits Lifelines2's utility. A fully operational tool will need to handle these additional data types.

Lifelines2 is not directly connected to the hospital database. While this allows for more platform-independent development, it requires pre-processing of WHC data. The process of obtaining and converting the data can take a long time because it takes collaboration among physicians and database administrators. We are currently working on a software middle layer that can significantly speed up the process, and reduce the amount of dependence on database administrators.

B. Main Benefits

There are three main benefits of Lifelines2. First, Lifelines2 has a far shorter learning time than SQL-like command-line languages. At the same time, Lifelines2 provides sufficient functionalities to many types of temporal medical queries. By lowering the barrier of using existing EHRs, we believe Lifelines2 increases the utility of EHR system to end-users, and broadens the pool of end-users.

Second, users can review and search faster and more accurately. As a consequence, better quality assurance metrics can be developed and policy impacts promptly measured, leading to better management of health care. In addition, this time saving can also allow physicians to better spend their time caring for patients, improving the quality of care.

Thirdly, the visual display of data and the interactive nature of Liflines2 facilitate discovery of messy data, which must be dealt with before analysis is performed or statistics are collected to avoid erroneous results.

The multiple case studies support our belief that appropriately chosen analytic operators (*e.g.* align, rank, filter, summarize) and user-controlled compact visualizations provide valuable tools for the study of large sets of patient records. Lifelines2 has been applied to case studies with tens of thousands of patients, covering hundreds of thousands of events, but we feel confident that we can deal with much larger sets of data.

The three specific benefits we claim stem from profound changes in the problem-solving processes we observed in ourselves and our physician partners. The analytic operators and visualizations provide a new language for thinking about and solving temporal sequence problems that has proven to be effective in medical and other domains. It takes new users a few minutes and a few examples to understand the relevant concepts, but then they begin to revisit old problems and develop new solutions. A further strong attraction is the visual presentation of results that quickly reveal data anomalies as well and surprising patterns that often lead to valuable insights. Just as Xrays and MRIs enable physicians to see patterns in human anatomy, visualizations of temporal event sequences may enable physicians to see patterns in patient histories.

Acknowledgements
This work was supported by Washington Hospital Center, MedStar Health and the grant RC1CA147489-02 from NIH-National Cancer Institute on "Interactive Exploration of Temporal Patterns in Electronic Health Records."

References

1. Shahar, Y, Dynamic temporal interpretation contexts for temporal abstraction. Annals of Mathematics and Artificial Intelligence, 1998;22:159-192.
2. Shahar, Y, Dimension of time in illness: an objective view, Annals of Internal Medicine, 2000; 132(1):45-53.
3. Augusto, JC, Temporal reasoning for decision support in medicine. Artificial Intelligence in Medicine. 2005;33(1):1-24.
4. Stacey, M, McGregor, C, Temporal abstraction in intelligent clinical data analysis: a survey. Artificial Intelligence in Medicine. 2007;39:1-24
5. Microsoft Amalga [Internet]. Microsoft. Available from: http://www.microsoft.com/amalga/
6. Murphy, S, Mendis, M, Hackett, K, Kuttan, R, Pan, W, Phillips, L, et al. Architecture of the open-source clinical research chart from informatics for integrating biology and the bedside. Proceedings of the American Medical Association Annual Symposium (AMIA '07). 2007:548-552.
7. Catarci, T, Costabile, MF, Levialdi, S, Batini C. Visual query systems for databases: a survey. Journal of Visual Languages and Computing, 1997;8:215-260.
8. Cheng, C, Shahar, Y, Puerta, A, Stites, D. Navigation and visualization of abstractions of time-oriented clinical data. Technical Report. Stanford University Section on Medical Informatics. 1997. Report No.:SMI-97-0688.
9. Cheng C, Shahar, Y. Intelligent visualization and exploration of time-oriented clinical data. Topics in Health Information Management. 1999;20:15-31.
10. Jensen, C S, Snodgrass, R T. Temporal data management. IEEE Transactions on Knowledge and Data Engineering. 1999;11:36-44.
11. Tansel, A , Arkun, ME, Ösoyoğlu, G. Time-by-example query language for historical databases. IEEE Trans. Software Engineering. 1989;15(4):464-478.
12. Hibino, S, Rundensteiner, EA. User interface evaluation of a direct manipulation temporal visual query language. Proceedings of the 5th ACM International Conference on Multimedia. New York, NY, USA, 1997:99-107.
13. Bruce, BC. A model for temporal references and its application in a question answer program. Artificial Intelligence. 1972;3:1-25.
14. Allen, JF. Maintaining knowledge about temporal intervals. Communications of the ACM. 1983;26:832-843.
15. Kahn, KM, Gorry, AG. Mechanizing temporal knowledge. Artificial Intelligence, 1977;9:87-108.
16. Zloof, MM. Query-by-example: A database language. IBM Systems Journal, 1977.
17. Kouramajian, V, Gertz, M. A graphical query language for temporal databases. Proceedings of Object-Oriented and Entity Relationship Modeling. Springer-Verlag. 1995:388-399.
18. Silva, SF, Schiel, U, Catarci, T. Visual query operators for temporal databases. Proceedings of the International Workshop on Temporal Representation and Reasoning, 1997:46-53.
19. Dionisio JDN, Cardenas, AF. MQuery: a visual query language for multimedia timeline and simulation data. Journal of Visual Languages and Computing. 1996;7:377-401.
20. Post, AR, Harrison, JH. Protempa: a method for specifying and identifying temporal sequences in retrospective data for patient selection. Journal of American Informatics Association. 2007;14(5):674-683.
21. Chittaro, L, Combi, C. Visualizing queries on databases of temporal histories: new metaphors and their evaluation. Data and Knowledge Engineering. 2003; 44:239-264.
22. Hochheiser, H, Shneiderman, B. Dynamic query tools for time series data sets: timebox widgets for interactive exploration. Information Visualization. 2004;3(1):1-18.
23. Ankerst, M, Jones, DH, Kao, A, Wang, C. DataJewel: integrating visualization with temporal data mining. Visual Data Mining. 2008:312-330.

24. Plaisant, C, Mushlin, R, Snyder, A, Li, J, Heller, D, Shneiderman, B, LifeLines: using visualization to enhance navigation and analysis of patient records. American Medical Informatics Association.1998 Annual Fall Symposium, Bethesda MD. 1998:76-80.
25. Bade, R, Schlechtweg, S, Miksch, S. Connecting time-oriented data and information to a coherent interactive visualization. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2004), ACM, 2004:105-112.
26. Plaisant, C, Lam, S, Shneiderman, B, Smith, M, Roseman, D, Marchand, G, et al. Searching electronic health records for temporal patterns in patient histories: a case study with Microsoft Amalga. Proceedings of the AMIA Annual Symposium. 2008:601-605.
27. Fails, J, Karlson, A, Shahamat, L, Shneiderman, B. A visual interface for multivariate temporal data: finding patterns of events over time. Proceedings of the IEEE Symposium on Visual Analytics Science and Technology. 2006:167-174.
28. Klimov, D, Shahar, Y, Taieb-Maimon, M. Intelligent selection and retrieval of multiple time-oriented records. Journal of Intelligent Information Systems [Internet]. 2009. Available from: http://www.springerlink.com/content/u5568658x7h15034/
29. Klimov, D, Shahar, Y, Taieb-Maimon, M. Intelligent visualization and exploration of time-oriented data of multiple patients. Artificial Intelligence in Medicine. 2010;49(1):11-31.
30. Wongsuphasawat, K, Shneiderman, B. Finding comparable temporal categorical records: a smiliarity measure with an interactive visualization. Proceedings of the IEEE Symposium on Visual Analytics Science and Technology. 2009:27-34.
31. Wang, TD, Plaisant, C, Quinn, A, Stanchak, R, Shneiderman, B, Murphy, S, Aligning temporal data by sentinel events: discovering patterns in electronic health records, Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI2008). ACM. 2008:457-466.
32. Wang, TD, Plaisant, C, Shneiderman, B, Spring, N, Roseman, D, Marchand, G, et al. Temporal summaries: supporting temporal categorical searching, aggregation, and comparison, IEEE Transactions on Visualization and Computer Graphics. 2009;15(6):1049-1056.
33. Wang, TD. Interactive Visualization Techniques for Searching Temporal Categorical Data [Ph.D. Dissertation]. Department of Computer Science, University of Maryland at College Park; May, 2010: Available from: http://hdl.handle.net/1903/10382.
34. Bederson, BB, Grosjean, J, Meyer, J. Toolkit design for interactive structured graphics. IEEE Transactions on Software Engineering. 2004;30(8):535-546.
35. Wongsuphasawat, K, Plaisant, C, Shneiderman, B, Querying timestamped event sequences by exact search or similarity-based search: design and empirical evaluation. Technical Report. Human-Computer Interaction Lab, University of Maryland. 2009. Report No.: HCIL-2009-20.
36. Janszky I., Ljung R. Shifts to and from daylight saving time and incidence of myocardial infarction. New England Journal of Medicine, 2008;359(18):1966-1968.