Foundations and Trends® in
Human–Computer Interaction
Vol. 5, No. 3 (2011) 207–298

DOI: 10.1561/1100000039

Interactive Information Visualization to Explore and Query Electronic Health Records

By Alexander Rind, Taowei David Wang,
Wolfgang Aigner, Silvia Miksch,
Krist Wongsuphasawat, Catherine Plaisant,
and Ben Shneiderman

Contents

Foundations and Trends® in
Human–Computer Interaction
Vol. 5, No. 3 (2011) 207–298

DOI: 10.1561/1100000039

the essence of knowledge

Interactive Information Visualization to Explore and Query Electronic Health Records

Alexander Rind[1], Taowei David Wang[2], Wolfgang Aigner[3], Silvia Miksch[4], Krist Wongsuphasawat[5], Catherine Plaisant[6], and Ben Shneiderman[7]

[1] *Vienna University of Technology, Austria, rind@ifs.tuwien.ac.at*
[2] *University of Maryland, USA, tdwang@partners.org*
[3] *Vienna University of Technology, Austria, aigner@ifs.tuwien.ac.at*
[4] *Vienna University of Technology, Austria, miksch@ifs.tuwien.ac.at*
[5] *University of Maryland, USA, kristw@cs.umd.edu*
[6] *University of Maryland, USA, plaisant@cs.umd.edu*
[7] *University of Maryland, USA, ben@cs.umd.edu*

Abstract

Physicians are confronted with increasingly complex patient histories based on which they must make life-critical treatment decisions. At the same time, clinical researchers are eager to study the growing databases of patient histories to detect unknown patterns, ensure quality control, and discover surprising outcomes. Designers of Electronic Health Record systems (EHRs) have great potential to apply innovative visual methods to support clinical decision-making and research. This work surveys the state-of-the-art of information visualization systems for exploring and querying EHRs, as described in the scientific literature.

We examine how systems differ in their features and highlight how these differences are related to their design and the medical scenarios they tackle. The systems are compared on a set of criteria: (1) data types covered, (2) multivariate analysis support, (3) number of patient records used (one or multiple), and (4) user intents addressed. Based on our survey and evidence gained from evaluation studies, we believe that effective information visualization can facilitate analysis of EHRs for patient treatment and clinical research. Thus, we encourage the information visualization community to study the application of their systems in health care. Our monograph is written for both scientific researchers and designers of future user interfaces for EHRs. We hope it will help them understand this vital domain and appreciate the features and virtues of existing systems, so they can create still more advanced systems. We identify potential future research topics in interactive support for data abstraction, in systems for intermittent users, such as patients, and in more detailed evaluations.

1

Introduction

Medical decision-making is a complex process. A patient's well-being depends on correct diagnosis and appropriate treatment. Physicians must incorporate large amounts of information such as a patient's status, symptoms, medical history, past and ongoing treatments, which are encompassed in the electronic health record (EHR). In addition, these records are an invaluable data source for clinical research and improvement of clinical quality, as they provide longitudinal health information about patient populations [49, 131, 138].

In recent years, many health care institutions have introduced EHR systems to replace their paper-based health records. However, current clinical information systems have focused on faster and cheaper management, storage, and sharing of EHRs. Unfortunately, EHR systems have been shown to have little positive effects on the quality of care, and in some cases have decreased quality [66]. A 2009 report by a committee of the National Research Council of the National Academies found that care providers spend considerable time entering data into EHRs for billing and legal purposes, but that this data rarely improves the quality of care, largely because EHR systems fail to provide cognitive support to healthcare providers, patients, and families [134].

Information visualization has the potential to address those issues and deliver much-needed cognitive support. Indeed, a 2012 report of the US Institute of Medicine [72], which focuses on improving patient safety, recommends "cross-disciplinary research" on "user-centered design and human factors applied to health IT." The report also notes that "Information visualization is not as advanced in parts of clinical medicine as compared with other scientific disciplines."

In the scientific literature, several information visualization techniques have been proposed that encourage users to explore EHR data visually, gain insights, and form hypotheses. Those systems have demonstrated some level of success, but it is difficult to get an overview and compare them. In this work we report on an extensive literature survey of visualization and interaction techniques applied to EHRs. We review and compare state-of-the-art research systems and examine their support for medical care, clinical research, and quality control. The focus is on information visualization techniques as opposed to medical imaging techniques. It also excludes techniques aiming to support the management of administrative or financial data.
This work presents:

(1) A survey of state-of-the-art information visualization systems from academic literature.
(2) A review of the visualization and interaction techniques found in 14 of these systems (Table 1.1) including strengths and weaknesses. These systems are categorized by the tasks and data (type, complexity, and scale) they support. Furthermore, there are compact descriptions of 32 additional EHR visualization systems.
(3) A summary of evaluation studies conducted in medical context.
(4) An overview of data visualization in commercial EHR systems.
(5) Recommendations and future research directions for information visualization in EHR systems.

Our analysis of single patient and multiple patient systems is written for both scientific researchers and designers of future user interfaces

Table 1.1. Overview of the 14 systems reviewed in detail.

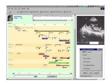	**LifeLines** (see Figure 4.1) University of Maryland [110]
	**MIVA** (see Figure 4.3) Indiana University [56]
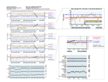	**WBIVS** (see Figure 4.4) University of Minnesota [107]
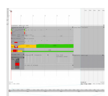	**Midgaard** (see Figure 4.5) Otto-von-Guericke University of Magdeburg [32]
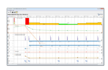	**VisuExplore** (see Figure 4.7) Vienna University of Technology [123]
	**VIE-VISU** (see Figure 4.9) University of Vienna [69]
	**Lifelines2** (see Figure 4.12) University of Maryland [147, 148]
	**Similan** (see Figure 4.13) University of Maryland [155]
	**PatternFinder** (see Figure 4.16) University of Maryland [53]
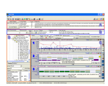	**VISITORS** (see Figure 4.18) Ben-Gurion University of the Negev [80, 81]
	**Caregiver** (see Figure 4.21) Fachhochschule Nordwestschweiz [36]
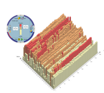	**IPBC** (see Figure 4.22) University of Udine [45]
	**Gravi++** (see Figure 4.24) Vienna University of Technology [67]
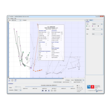	**TimeRider** (see Figure 4.25) Danube University Krems [122]

for EHR data analysis. These interface designers face a substantial challenge in understanding medical care, clinical research, and quality control sufficiently well to create effective interfaces. If these designers appreciate the features and virtues of existing systems, they will be more capable in creating still more advanced systems.

We first provide background information on information visualization in the medical domain, highlight its significance, and compare this survey to existing work. The *Methods* section presents our approach to searching relevant literature and our review criteria. The *Results* section presents 14 information visualization systems and briefly describes related systems. The *Discussion* section evaluates the 14 systems using our review criteria, reports on evaluation studies, gives an overview of commercial systems, explains limitations, and provides recommendations for future work.

2

Background

Medical data is often large, complex, disorganized [131], and may reside in separate databases or in data warehouses. These characteristics make medical information difficult to integrate and analyze. However, critical analysis of medical data can improve health care and have a positive impact on lives. An EHR can be characterized as "the complete set of information that resides in electronic form and is related to the past, present and future health status or health care provided to a subject of care" [113, p. 104]. As EHR systems become more prevalent, the need to have effective support to access and understand them becomes urgent. Information visualization is one way to improve the understanding of complex data and consequently increase the value of electronically available medical data [42].

First, we describe the challenges of analyzing medical information and how interactive visualization systems can help users overcome these challenges.

2.1 The Challenges

To improve care EHR data is used in two basic ways: (1) physicians need to understand a particular patient's data to make medical decisions; and (2) clinical researchers and quality improvement analysts need to understand patient population's data to help establish new clinical knowledge or assess the overall quality of care. While reliable data entry and retrieval is the basic function of EHR systems, powerful exploration methods and rich query capabilities are needed to realize the benefit of EHR systems and clinical data warehousing. Unfortunately, methods to explore, query, and interact with the data have not improved to deal with the size and complexity of the data available to date. As a result, data in these systems are often relegated to the usage pattern of "Write-Once, Read Never" [114]. Efforts to improve EHR accessibility to physicians have met with varying degrees of success. Christensen and Grimsmo [46] reported that EHR systems used by general practitioners improved the availability of patient information, but finding information within each record was not easy enough: 37% of study respondents reported that they sometimes gave up searching for information because it was too time consuming [46]. EHR systems facilitate record-keeping and management of data, but can introduce new medication errors through poor interaction design [31, 82].

Without proper interaction design, systems can pose high technical barriers to their end-users. For example, within a single patient record the temporal information is typically fragmented over multiple screens and multiple tables, making it difficult to review the sequence of events [87, 107]. Another example is an EHR data warehouse system that requires its users to use command-line style queries to perform searches. Command-line query languages such as SQL or its extensions for the medical domain [75, 99, 103] are very difficult to learn and temporal queries are extremely hard to specify. Some systems encapsulate query languages with graphical user interfaces that are easier to learn. However, these systems often constrain the expressiveness of query languages in order to keep the user interface manageable. As a result, even state-of-the-art systems rarely support rich temporal queries [97, 109] despite the need for temporal information in clinical research.

2.2 Information Visualization

"Computer-based visualization systems" can be defined as "provid[ing] visual representations of datasets intended to help people carry out some task more effectively" [96]. Information visualization techniques focus on datasets with nonspatial data attributes and discrete observations [141]. Scientific visualization, on the other hand, concentrates on visualizing real objects with spatial dimensions—typically three-dimensional (3D), for example, tumor nodes or blood vessels (cp. [116]). Since the EHR encompasses data relevant for both information visualization and scientific visualization, EHR systems often combine different visualization methods. Information visualization relies on suitable mappings of abstract data to compact representations to convey meaningful information quickly, for example a sparkline of glucose readings [143]. Thus, information visualization is suitable for exploring and querying the heterogeneous and temporal data found in EHRs. Furthermore, interactive information visualization is one way to enable exploratory analysis, an important partner to statistical confirmatory analysis [144].

Interaction is "at the heart" of information visualization [132]. Interaction allows users to dynamically change the mapping of the data (e.g., color, shape, size), the view of the mapping (e.g., zoom, pan, rank), or the scope of the data being visualized (e.g., search, filter). Aggregation or clustering can also be supported. Information visualization research aims to combine the processing power of modern computers with human cognition and visual abilities to better support analysis tasks.

The visual information-seeking mantra "Overview, Zoom and Filter, then Details on Demand" [129] provides a simple guideline to deal with massive, disorganized medical data. Users can grasp distributions, trends, and anomalies once large amounts of data have been compressed into concise overviews. By using effective strategies of zoom and filter, users can selectively focus on interesting data points. Finally, details can be shown when users request them. Information visualization researchers have proposed many techniques to deal with large amounts of multidimensional data (e.g., VisDB [77], TableLens [118]), effective and learnable search strategies (e.g., dynamic query [22], Attribute Explorer [145]), and dynamic visual rearrangements that improve visual exploration (e.g., Fisheye [60], semantic zoom [65], Line Graph Explorer [79]). However, depending on the characteristics of the data and the user tasks, different techniques should be used, and the primary goal for this survey is to compare and contrast the currently available approaches.

Visualization of temporal data. The EHR contains a patient's "past, present, and future" [113] medical data. Since the health state of a patient changes over time, especially in the course of medical treatment, time and temporal data play an important role in exploring and querying EHRs. However, temporal data has special characteristics that distinguish it from other dimensions (e.g., intervals, calendar structure, and cyclic events). Visualization methods that consider these characteristics can allow more effective analysis of such data [25, 26].

A good coverage on visualization of temporal data is given in the book by Aigner et al. [26] who present a survey of 101 visualizations, which includes some EHR visualization systems. Although, other visualization methods for temporal data such as Continuum [29], GROOVE [85], LiveRAC [93], or TimeSearcher [38, 39, 68] can in principle be applied to visualize EHRs based on the underlying data characteristics, we put a focus on methods that have been developed, applied, and/or tested in medical contexts. Thus, the presented methods have shown their suitability for the specific application case with its specific requirements already.

2.3 Related Surveys

There are several publications on information visualization for medical applications. However, these publications are scattered across the venues of several different target audiences (visualization, human-computer interaction, medical informatics, medicine). We know only two venues that focus on EHR visualization: a series of workshops organized by the Human-Computer Interaction Lab at the University of Maryland [6], and the recent VisWeek workshop series "Visual Analytics in Healthcare" [21].

Table 2.1. Coverage of reviewed systems in related work.

	year of publication	*LifeLines*	*MIVA*	*WBIVS*	*Midgaard*	*VisuExplore*	*VIE-VISU*	*Lifelines2*	*Similan*	*PatternFinder*	*VISITORS*	*CareGiver*	*IPBC*	*Gravi++*	*TimeRider*
Chittaro [42]	2001	●				○					○				
Kosara and Miksch [83]	2002	●	○				●								
Lungu and Xu [89]	2007	●	○			○									
Aigner et al. [23]	2008	●	○		●	○	●				○		●	●	
Combi et al. [49]	2010	●				○					●		●		
Lesselroth and Pieczkiewicz [86]	2011	●	○	●		○	●	●			●				

●: system reviewed (based on reference list).
○: similar system reviewed (e.g., KNAVE → VISITORS, Graphical Summary → MIVA).

Six other surveys provide partial coverage of the topic (Table 2.1): Chittaro's article [42] introduces the visualization field to the artificial intelligence in medicine community. It presents a "gallery", that is, a panorama of relevant visualization systems, which is loosely structured, comparatively short, and includes only few systems concerned with EHRs. Kosara and Miksch's review [83] does not examine concrete visualization systems but visual representation techniques. They focus on three tasks: visualizing measured data, incidents and symptoms, and treatment planning. For each task they establish requirements and assess the techniques based on these requirements. For example, the requirements for visualization of measured data are intuitiveness, focus+context time, focus+context data, combination of values, seeing developments, finding patterns, and discovering intervals. The techniques for visualization of measured data are line plot, Graphical summary of patient status, VIE-VISU, and spirals. Overall, this review makes an interesting counterpart to the work at hand. Lungu and Xu's book chapter [89] covers visualization systems for many areas of biomedical research. Concerning medical records, it describes only two systems for analysis of EHR data (LifeLines, The Cube) and two systems for management of treatment plans. The book chapter by Aigner et al. [23] studies visual methods for management of clinical guidelines. As a subpart of this survey they present six EHR visualization systems

(LifeLines, Midgaard, VIE-VISU, IPBC, Gravi++, KNAVE II) and systems that combine clinical guidelines with EHR data (e.g., CareVis). The book by Combi et al. [49] closes with a chapter on the visualization of temporal clinical data and knowledge. Here, they describe five visualization systems (IPBC, KHOSPAD, KNAVE II, Paint Strips, VISITORS) in detail and characterize them along four dimensions: subject cardinality (single/multiple patients), concept cardinality (single/multiple variables), abstraction level (raw data, abstract concepts, knowledge), and temporal granularity (single, single but variable, multiple). Lesselroth and Pieczkiewicz's book chapter [86] surveys different visualization strategies for EHRs. This work provides pointers to a large number of visualization systems, but they are not systematically described. Its narrative is structured along the sections multimedia, smart dashboards to improve situational awareness, longitudinal and problem-oriented views to tell clinical narratives, iconography and context links to support just-in-time information, and probability analysis and decision heuristics to support decision analysis and bias identification.

Our survey differs from the previous work in that it focuses exclusively on visualization systems for exploring and querying EHR data. It covers significantly more systems of that kind than other surveys and it is systematically structured using review criteria that designers must consider when creating appropriate information visualization systems.

3

Methods

This work studies the application of information visualization techniques to patient data in EHRs for medical care and clinical research. Its content is based on a survey of the academic literature. Here we describe the search strategy and the inclusion criteria for this survey. Then we introduce the review criteria, which help us to structure our work.

3.1 Literature Search

We collected research papers describing the design, implementation, or evaluation of information visualization systems. Our search strategy involved three steps. First, we searched in four electronic databases (ACM Digital Library, IEEE Xplore, Google Scholar, PubMed) and in the archives of two medical informatics journals (International Journal of Medical Informatics, Journal of the American Medical Informatics Association). Our principal search terms were "visualization of patient record". Another fruitful approach we adopted was to search for papers citing early articles such as Powsner and Tufte's graphical summary of patient status [114] and LifeLines [110] in the databases mentioned above. Next, we manually scanned the reference lists of the most relevant papers we had already found. Finally, we manually identified additional papers in conference proceedings (e.g., HCIL Workshop 2008, IEEE Conference on Visual Analytics Science and Technology (VAST), International Conference on Information Visualisation (IV), IEEE VisWeek Workshop on Visual Analytics in Health Care) and journals (e.g., Information Visualization Journal, IEEE Transactions on Visualization and Computer Graphics) of the information visualization community, which were published in 2007 and later, so that we could include state-of-the-art research that is not yet cited elsewhere. We also contacted the authors of included work and asked for more recent publications.

The decision whether a system should be included in this work was based on a number of criteria:

- The EHR consists of many kinds of data and spans entire patient histories. We are interested in systems that work with a respectable number of data items per patients and take temporal evolution of patient data into account. We excluded systems that show only a snapshot of patient state or few extracted items.
- Further, information visualization deals with discrete non-spatial data such as diagnoses, test orders and results (e.g., blood tests and normal/abnormal ratings), heart rate, or drugs orders. This does not include the visualization of spatial data, such as X-ray images, computed tomography data, or other medical images, for which the methods of scientific visualization are more relevant [116]. Display of free text (e.g., the full text of notes or discharge letters) is also out of this review's scope, even though data extracted from such texts can be explored with such tools (e.g., [139]). However, systems that combine information visualization with medical images or text are included.
- Interaction is intrinsic to information visualization, therefore this survey includes only interactive applications and skips approaches that produce only static pictures. Nevertheless,

various interaction techniques exist and even simple panning or geometric zooming is considered.
- The application domain of the work surveyed here includes patient care, clinical research, and quality control. It excludes the visualization of medical guidelines and information visualization systems designed for administrative tasks with or in relation to patient data (e.g., billing, scheduling).

Besides these explicit inclusion criteria, our search strategy had an implicit effect on the survey systems that should be considered here:

- A search of scientific literature unfortunately excludes many existing commercial systems. On the one hand, these systems are very relevant to our intended audience. However, on the other, web resources on these systems do not discuss their visualization design features, evaluation methods, or results as extensively and systematically as an academic paper would, making them poor subjects for our review criteria (see next section). Instead, we present an overview and a few examples of commercial systems (Section 5.3 *Patient Data Visualization in Commercial EHR Systems*) so that readers can draw connections to the systems in this survey.
- We survey systems that have been applied in the health care domain and not other systems that, in principle, could be applied there.
- Finally, we look at visualization systems and not on individual visualization or interaction techniques.

After initial research and preliminary coding of the review criteria, we selected the most original or most relevant systems, which we describe in more detail. Other systems that are similar to earlier work but with limited novelty in visualization or interaction are cited and briefly explained throughout the survey. This survey covers all identified systems that fulfilled our inclusion criteria, but of course there may be valuable ideas in unpublished systems or in papers we have missed.

3.2 Review Criteria

After an initial review of the systems and discussion among the authors we selected four review criteria that would help structure the results and discussion sections. Those criteria are data type covered, support for multiple variables, support for visualizing data from one versus many patients, and support for different user intents.

3.2.1 Data Types Covered

Visualization systems vary greatly as a function of the data type(s) they can handle. On a domain-independent level, systems can handle categorical (includes nominal or ordinal scale) or numerical variables (interval or ratio scale), or both (cp. [136]). Categorical variables have values from an unordered or ordered set. Nominal (unordered) examples include diagnoses (e.g., flu, lupus, cancer) or interventions (appendectomy, cesarean section, blood transfusion). Examples of ordinal (ordered) data type include cholesterol levels ("low", "normal", "high") and severity of a symptom on a scale from "none" to "severe". Numerical variables can have an interval or ratio scale. Examples are temperature in degrees Celsius (interval scale) or creatinine value in mg/dL (ratio scale). It is also possible to describe the scale of numerical variables as continuous or discrete (real or integer numbers). Some information visualization systems can only deal with a single data type while others can combine multiple types.

On a medical-domain level systems or techniques may be adapted to certain data types such as medical tests (e.g., blood glucose), diagnoses (e.g., diabetes), or treatments (e.g., medication or surgery).

3.2.2 Support for Multiple Variables

Physicians often need to analyze the codevelopment of two or more variables (e.g., a drug administration and test results). Different systems represent multiple variables differently, and we catalog these visualization and interaction techniques. We also look at the number of variables described in the original publications to propose an estimate of the appropriate number of variables the technique can handle.

3.2.3 Support for One versus Multiple Patient Records

A strongly differentiating criterion is whether the information visualization system supports the exploration and querying of a single patient record or collections of patient records. Scenarios that require working with multiple patient records include monitoring patient cohorts [36], quality assurance [45], alarm specification [109], clinical trial recruitment, and observational research using existing patient data [147]. We examine how users can query multiple EHRs but also how users can interactively explore individual patients within the result set.

3.2.4 Support for User Intents

Finally, information visualization systems can be compared by their interaction features. As stated above, interaction is an integral part of information visualization and contributes a large part of the benefits of EHR visualization systems. A wide choice of interaction techniques is implemented by the systems in this survey and, often, interaction techniques are interwoven with the visual representations. Therefore, we use the user intent model proposed by Yi et al. [157]. This categorization focuses on "what a user wants to achieve" and describes an interaction technique by one of seven user intents (see below). In order to better describe EHR visualization systems, we extended the model with 20 subintent categories:

(1) **Select:** Mark a subset of the dataset as interesting.
   - to keep track of selected items for a short term while the visualization is changed
   - to manage groups, for example adding or removing patients to groups

(2) **Explore:** Show a different part of the dataset.
   - to navigate in time (e.g., panning and zooming the time axis)
   - to add or remove parameters to the visualization
   - to add or remove patients to the visualization

(3) **Reconfigure:** Rearrange the visual layout of the dataset.

- by repositioning items manually (freely or by some constraints)
- by sorting items along an axis
- by other adjustments of an axis (e.g., alignment to a relative timescale, distortion to see some items in focus and some in context)
- by applying another technique to avoid occlusion (e.g., 3D camera movement)

(4) **Encode:** Change the way each item of the dataset is represented.

- by switching to a different visualization technique or opening it in a new view
- by varying visual encoding (e.g., map outcome to item color, encode severity as item size)

(5) **Abstract/Elaborate:** Show less or more detail.

- by abstraction of one or more parameters (e.g., abstraction of a series of temperature readings over 37.5°C into a period of fever or subsuming different medication as beta-blockers)
- by temporal data binning (e.g., aggregate parameter values either by fixed time intervals or for as long as they have the same value)
- by showing details of items (e.g., in a tooltip)

(6) **Filter:** Show or highlight something conditionally.

- by patient status without considering time or development over time
- by development over time like event sequences (e.g., surgery after stroke) or value trends (e.g., increasing cholesterol) without time constraints
- by time constraints (e.g., relapse within 3 weeks after discharge, surgery in May 2009)

(7) **Connect:** Show related data.

- to show patient/patient group relationship
- to brush items in other representations
- to brush items for other variables at the same point of time or of the same patient

While most systems support each of these intents, some are more elegant than others. We describe the features in each system and the intents they support. We hope that understanding the trade-offs made by these systems will help design create more effective systems. Although we do not explicitly map each system feature into one or more intents in the system descriptions, we trust that readers can infer them from the detailed descriptions. Table 5.3 summarizes the strengths and weaknesses of systems in terms of user intents.

4

Results

In this section the most representative EHR visualization systems are described in detail, along with the mention of other related systems or variants. Because the criterion "Support for one versus multiple patient records" has the largest impact on system design, the survey is first organized using this criterion, starting with systems that deal with a single patient and followed by systems that analyze collections of patient records. Below that, we grouped the systems under four headlines each that illustrate common features relevant for the exploration and querying of EHRs. Yet, these groups cannot subsume all characteristics of the systems, for which the textual description and the discussion tables should be consulted.

4.1 Visualization of a Single Patient Record

Visualizations of a single patient record support core clinical care tasks such as providing an overview of the patient history to understand trends, identify significant events, and spot omissions in data or treatment.

4.1.1 Events over Time

Developed in the late 1990s, **LifeLines** [110] is a seminal work of the visualization of personal medical histories. It uses a set of line segments to represent the events and episodes in a single patient record (Figure 4.1). The focus is on categorical data, using color to show normal or abnormal states. LifeLines supports the visualization of many medical variables on the same screen (e.g., 32 distinct types of events can be found in Figure 4.1). The line segments are distributed along the horizontal time axis, which can be zoomed and panned to reveal more or less details. Today's date is marked with a thick vertical line. Future events can be seen on the right of that vertical line (e.g., appointments or the scheduled end of a treatment). Along the vertical axis the data is grouped by facets that represent different aspects of the record such as "problems", "diagnoses", "tests", and

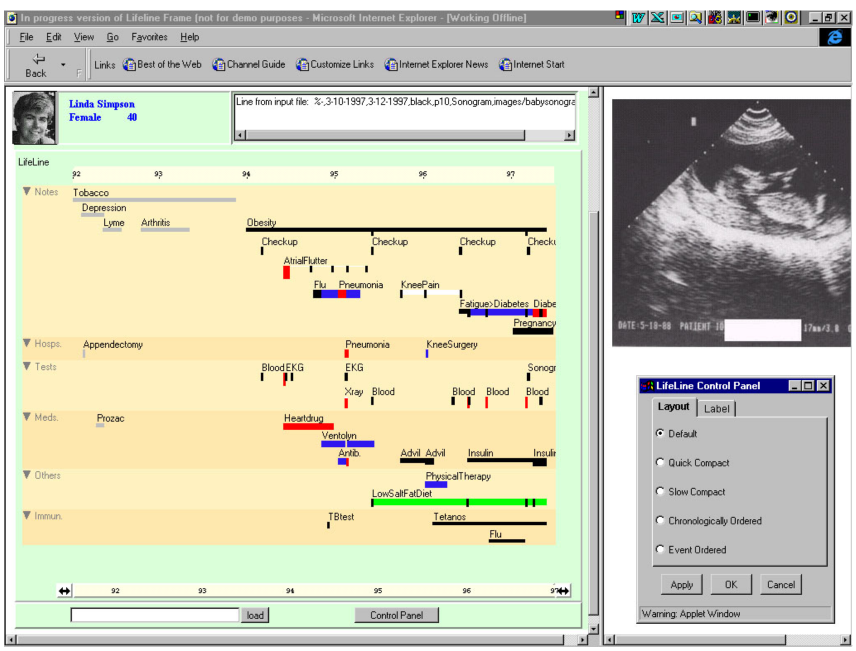

Fig. 4.1 LifeLines [110]: time line visualization of personal medical histories. Horizontal lines represent the events and episodes in a single patient record. Line color, height, and captions are used to convey information. A detail panel to the right shows further information about a selected item such as an ultrasound image. *Image by the authors.*

"medications". Facets can be closed into a much thinner colored representation called a silhouette that merely signifies the absence or presence of data. LifeLines groups similar events such as orders of the same drug in the same vertical position and allows aggregation of sets of events onto summary events (e.g., all consecutive orders of the same drug or class of drug can be aggregated).

To access more details, users can move the mouse cursor over a line and read a long label. Or they can double-click on any line and open documents or images in the right side, turning the visualization into large menu to access details. A string search function allows users to highlight events that have the search term in their description. LifeLines has been an inspiration for many other visualization tools.

KHOSPAD [51] show medical events with support for temporal granularity and indeterminacy. For this, it provides a complex notation for events, which for example shows the minimum duration, and an additional view on temporal relationships between events. **AnamneVis** [158] and an earlier system [159] apply vertical lines or boxes for visualizing the "when" aspect of data in an EHR. In addition, these two systems provide other views such as node link diagrams for the "why" aspect, a body shape for the "where" aspect, and a sunburst visualization for the "what" aspect of a patient's diseases. Other examples using horizontal lines for medical data are **Patient History in Pocket (PHiP)** [30], which is optimized for low-resolution mobile devices in epilepsy treatment, and a visualization tool for the identification of patients with poor medication adherence [90].

4.1.2 Numerical Data over Time

Point plots are a common technique for visualization of numerical data. Powsner and Tufte's **graphical summary of patient status** [114, 115] (Figure 4.2) is a well known yet non-interactive visualization that demonstrates point plots and small multiples. It can present two dozen variables by using common axes and nonlinear scales. It compares numerical data to laboratory reference values or recommended drug doses and visually highlights abnormal data.

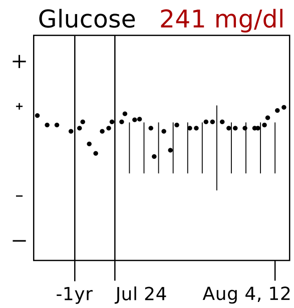

Fig. 4.2 Graphical summary of patient status [114, 115] is composed of two dozen point plots like this. The vertical axis is scaled based on five reference value ranges (critically elevated, elevated, normal, reduced, critically reduced) and abnormal values are visually highlighted. The horizontal axis is scaled nonlinearly to present the complete patient history with data of the current treatment (here, Jul 24 until Aug 4) being shown in more detail. The number in the top right corner shows the most recent value.
Image by the authors, based on the concept by Powsner and Tufte.

The **Medical Information Visualization Assistant (MIVA)** [54, 56], formerly known as Critical Care Patient Data Visualization system (CPDV) [55], presents many of the concept from Powsner and Tufte and include them in an interactive user interface (Figure 4.3). It uses point plots to visualize the change of numerical values over time. Each variable is displayed in a separate plot. A gray band in the background denotes the normal range of the variable in each plot. The point plots are stacked vertically and aligned on a shared time axis, which users can zoom and pan. Meanwhile, the small plots on the right always show the most recent developments and the current value is displayed as a large label. The variable label is color-coded to indicate recent abnormal values. Variables can be added from a list or reordered by drag and drop. Users can display all variable values at a certain time through the "scrubber", which is represented as a vertical red line. On top, there are two small panels with clinical text notes, clinical interventions, and other medical events. MIVA can thus also present some categorical data as contextual information.

Systems that plot numerical data over time are often found in commercial systems (see *Discussion* section) and several other systems were found in the literature survey: Rasmussen and Starren's

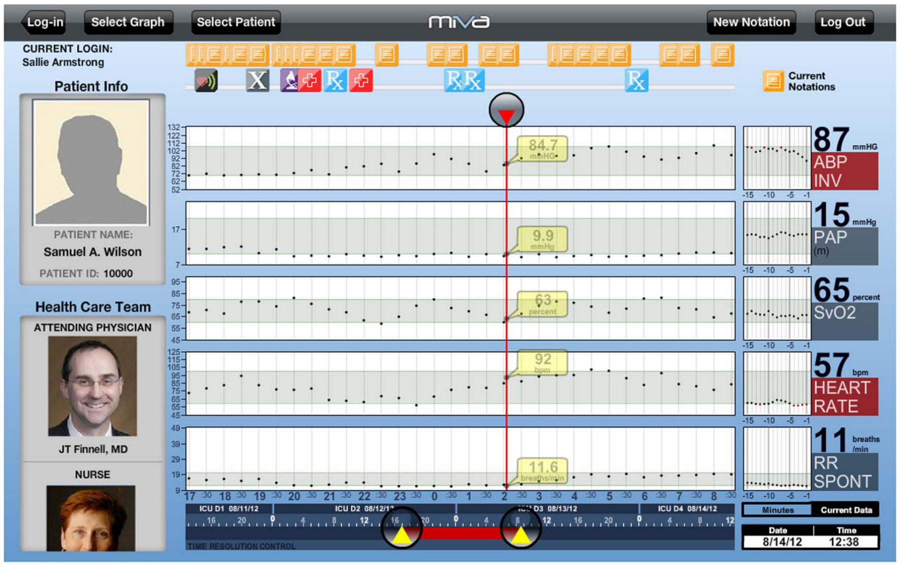

Fig. 4.3 Medical Information Visualization Assistant (MIVA) [54, 56]: set of point plots for intensive care data, which can be added and reordered. The gray bands in the background denote the normal range of the variable in the plot. The large labels and the small plots on the right show the current patient status and the most recent development. On top, the icons represent clinical text notes, interventions, and other events as aligned to the interactive time axis in the bottom. *Image courtesy of Anthony Faiola.*

Augmented Interactive Starfield Display [119] visualizes blood glucose values over time in a point plot. Fonseca et al. [59] present a visualization using overlaid line plots in context of a data integration project. **The Cube** [57] visualizes numerical data in a 3D parallel coordinate plot. Sparklines, line plots as small as a line of text [143], are evaluated by Bauer et al. [34] and reportedly used at Lucile Packard Children's Hospital at Stanford [104]. The **integrated graphical information display (IGID)** for intensive care patients evaluated by Anders et al. [28] is composed of vertically stacked line plots.

4.1.3 Heterogeneous Data on a Common Time Axis

Combined visualization of event and numerical data over time is a feature of many single patient record systems. These systems typically provide multiple charts of different representation techniques, which are aligned on a common horizontal time axis.

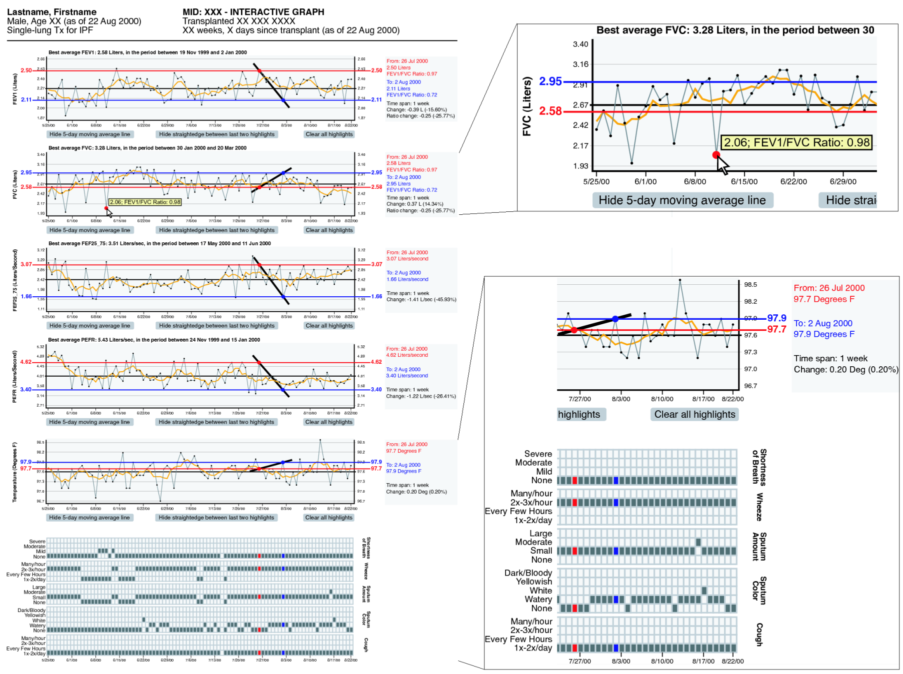

Fig. 4.4 Web-based Interactive Visualization System (WBIVS) [107] aligns line plots and matrix plots for lung transplant home monitoring. Parts of the screen shots are zoomed in on the right to give a better view of the details.
Image courtesy of David Pieczkiewicz, zoomed details added by the authors.

The **web-based interactive visualization system (WBIVS)** [107] for lung transplant home monitoring data shows both numerical and categorical pulmonary data (Figure 4.4). It combines line plots for multiple numerical variables and matrix plots for categorical variables. In total the application visualizes ten variables.

When users select a data point, the data points for the same time period are highlighted in all the other charts (with color and horizontal lines). Details about the last two selected points are shown on the right of the graph and a thick line is drawn between the two points to visualize the rate of change. If the mouse pointer moves over a point, a tooltip shows its value. The information provided on the right or in the tooltip includes derived values for lung function, absolute and relative change. A five-day moving average line can be displayed in order to highlight trends or hidden to reduce clutter.

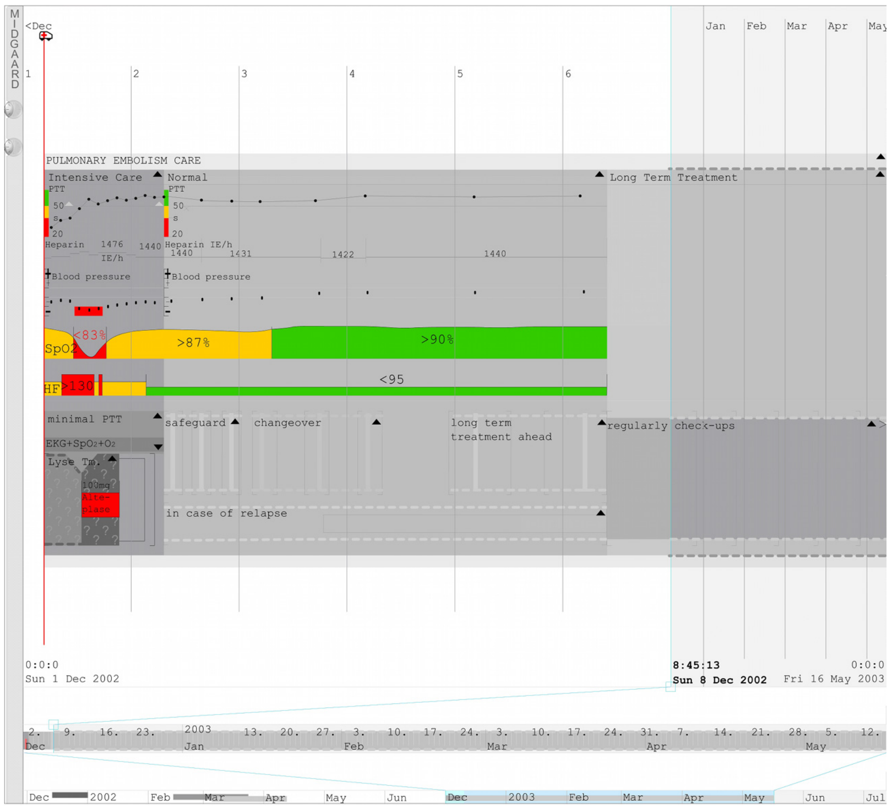

Fig. 4.5 Midgaard [32]: This visualization for intensive care includes a powerful semantic zoom, which allows users to smoothly zoom in and out and progressively reveal more details through animated transitions. *Image by the authors.*

Midgaard [32] is a visualization system designed for intensive care. It integrates the display of numerical data with graphical representations of medical treatment plans (Figure 4.5).

Midgaard provides a sophisticated semantic zoom visualization technique for numerical variables, which incorporates knowledge about the variable. Depending on the zoom level and the available screen area, Midgaard adapts the level of detail and shows the data either as (a–b) colored background, (c) colored bars, (d) area charts, or (e) augmented line charts (cp. Figure 4.6). For this, it calculates categorical abstractions of the numerical data. A time series of blood pressure measurements, for example, could be abstracted to periods of "normal

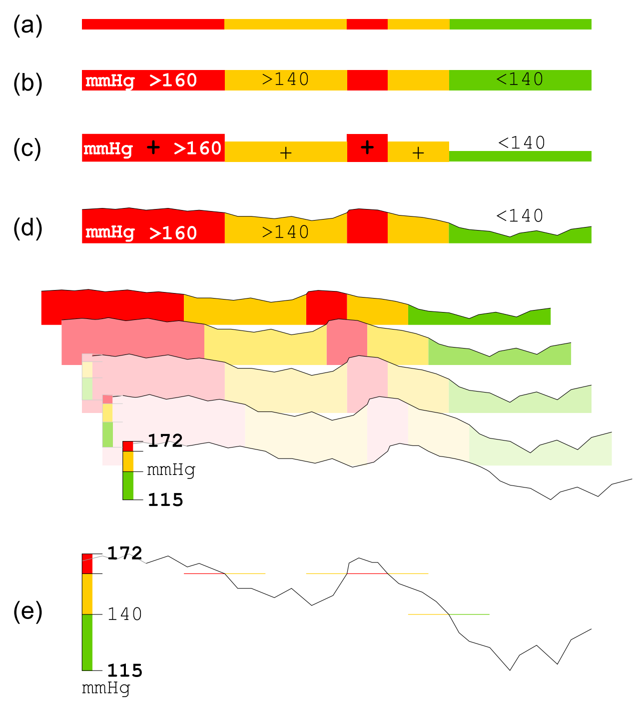

Fig. 4.6 Semantic zoom visualization technique for numerical variables in Midgaard [32] with five different levels of detail: (a) colored background, (b) colored background with labels, (c) colored bars, (d) colored area charts, and (e) augmented line charts.
Image by the authors.

blood pressure" (green), "increased blood pressure" (yellow), and "critical blood pressure" (red). This visualization technique allows Midgaard to show a multitude of variables in juxtaposition and aligned to a horizontal time axis. Users can resize the visualization of each variable and thus change the level of detail by dragging its border up or down. A controlled user study [27], which compared this technique to the visual representations in KNAVE II [128], showed that task performance was similar or better with Midgaard's technique and in particularly faster for more complex tasks.

Likewise, Midgaard's visualization technique adapts the level of detail along the time axis. If a user zooms in, individual data points will be drawn. Optionally, additional details of the data points such as measurement deviation, trustability, or valid time will be shown.

When users zoom out, Midgaard progressively switches to more compact graphical elements, for example, from point plots to lines to a high-frequency data display based on Tukey box-plots [144] or Information Mural [73].

In addition to numerical data, Midgaard visualizes medical treatment plans as colored bars, which can contain further bars representing subplans.

To navigate and zoom over time, users interact with two time axes that are located below the visualization area. The bottom time axis shows an overview of the full time span covered in the record. Users can select a visible time range for the middle time axis and the visualization area. On the middle time axis users can add distortion borders to the visible time range, making it possible to view some time periods in more detail than others, a visualization technique known as focus and context [48]. Users can use this control either by dragging borders on both time axes and the visualization area or by using a predefined layout. When the mouse hovers over a data point, its exact value appears in a tooltip.

In a similar fashion, **VisuExplore** [111, 123] shows patient data in multiple views that are aligned to a horizontal time axis (Figure 4.7). Yet, it is more flexible and allows for each view the selection of a visualization technique to display one or more variables. Well-known, deliberately simple visualization techniques are available by default in order to make the system easy to learn and intuitive to use: line plots and bar charts for numerical data, event charts and charts based on LifeLines [110] for categorical data. Moreover, VisuExplore is extensible and advanced visualization techniques based on the semantic zoom technique of Midgaard [32] and horizon graphs [120, 126] have already been added. VisuExplore provides rich interaction for visual exploration and flexible adaption. Users can freely resize and reorder views, pan and zoom the time axis, and show details in a tooltip or an optional data table view. In addition, a measure tool allows them to determine the time interval between any two points on the screen.

The **Time Line Browser** [52] is an early visualization prototype for heterogeneous diabetes data. **KNAVE** [127] and **KNAVE II** [91, 128]

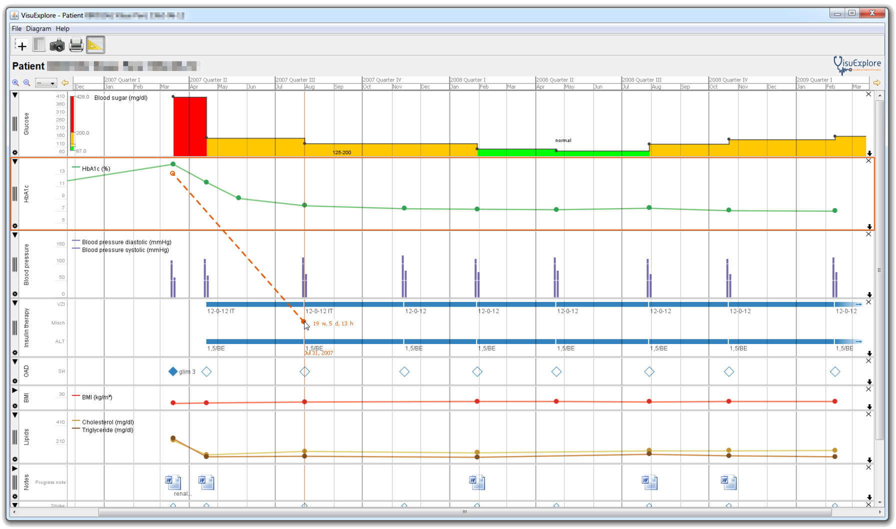

Fig. 4.7 VisuExplore [111, 123] providing overview of ten diabetes examinations with line plots, bar charts, event chart, and representation techniques based on LifeLines and Midgaard. A time span is measured starting from the peak HbA1c value (dashed orange line). *Image by the authors, used under CC-BY-ND license.*

use clinical knowledge bases and temporal data abstraction to visualize raw and summary data from an EHR (Figure 4.8). They are predecessor systems of the VISITORS system [80, 81], which is described in more detail below. **CareVis** [24] displays numerical patient data in juxtaposition with a graphical representation of the applied treatment plan. **Homecare** [43] combines several charts along a horizontal time axis on a low-resolution mobile device. Further systems in this group are **ICUFiles** [125], **TimeLine** [37], and an interactive real-time visualization environment for patients with heart failure [117].

4.1.4 Snapshots of Patient State

VIE-VISU [69] uses glyphs to represent a patient's status in neonatal intensive care. A series of glyphs shows changes over time (Figure 4.9). All glyphs use the same measures and scale, so they can be compared (cp. small multiples [142]). A glyph is a composite graphical object that has different geometric and visual attributes, which are used to encode multivariate data [78, 151]. VIE-VISU encodes categorical variables

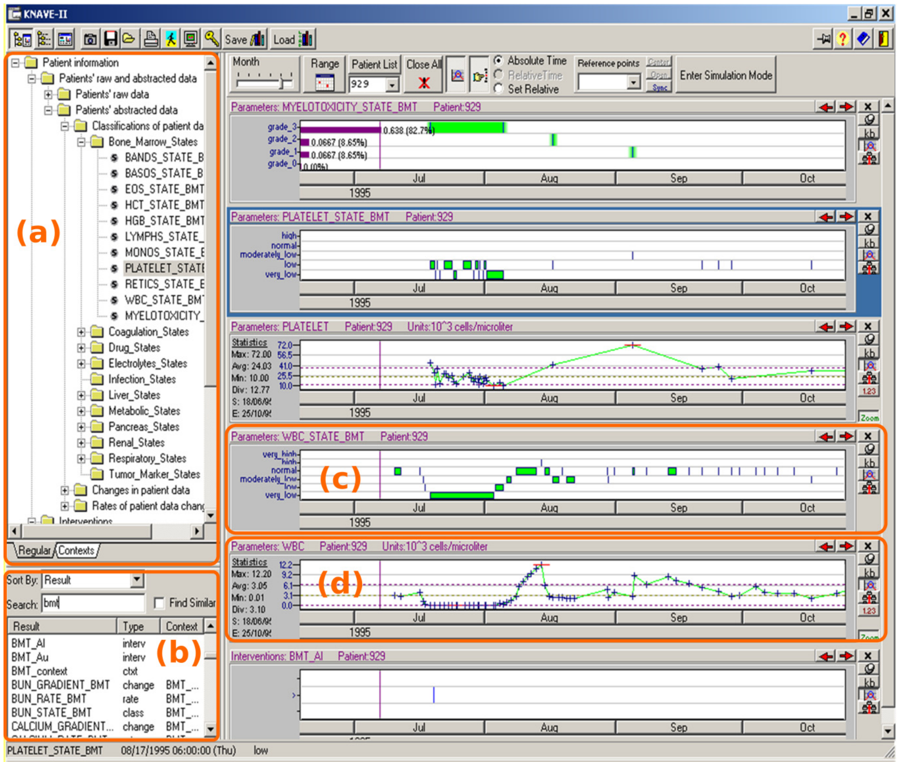

Fig. 4.8 KNAVE II [128] providing overview of an oncology patient's data: (a) on the left, a browser of the domain ontology allows the user to add abstract concepts and raw data variables for visualization. (b) In addition there is a string search for concepts. (c) Abstract concepts have an ordinal scale and duration. They are represented with line segments similar to LifeLines. (d) Line plots display numerical raw data.
Image courtesy of Yuval Shahar and Denis Klimov, markup added by the authors.

with geometrical shapes and their color, and encodes numerical variables by the size of glyph elements. In total 15 variables can be encoded in each glyph, covering important circulatory, respiratory, and fluid balance variables (Figure 4.10). Variables are grouped by physiological systems in the glyph, for example circulatory parameters are mapped to a triangle on top of the glyph. The width of the triangle's base shows the heart rate, the triangle's height represents the blood pressure, and the color displays catecholamine on a five level ordinal scale. By default VIE-VISU shows 24 glyphs, one per hour. The supported

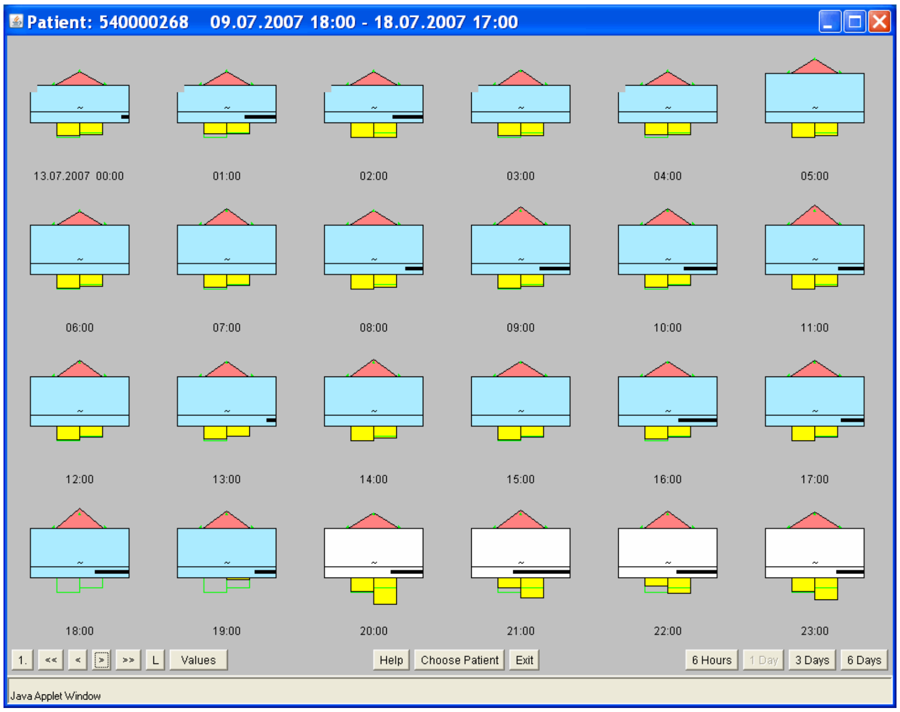

Fig. 4.9 VIE-VISU [69]: glyph-based, small multiple visualization for intensive care. A single glyph is a composite object that represents 15 different variables. Different visual variables like color or size are used and combined to form a meaningful whole.
Image courtesy of Werner Horn.

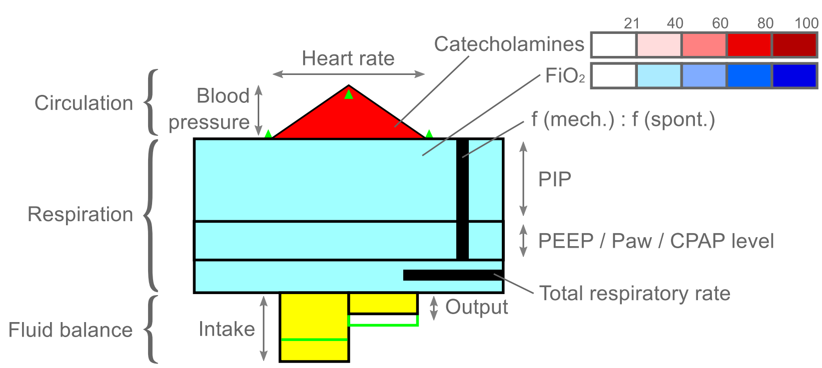

Fig. 4.10 Legend to the visual encoding of medical variables in VIE-VISU [69].
Image courtesy of Werner Horn, adapted by the authors.

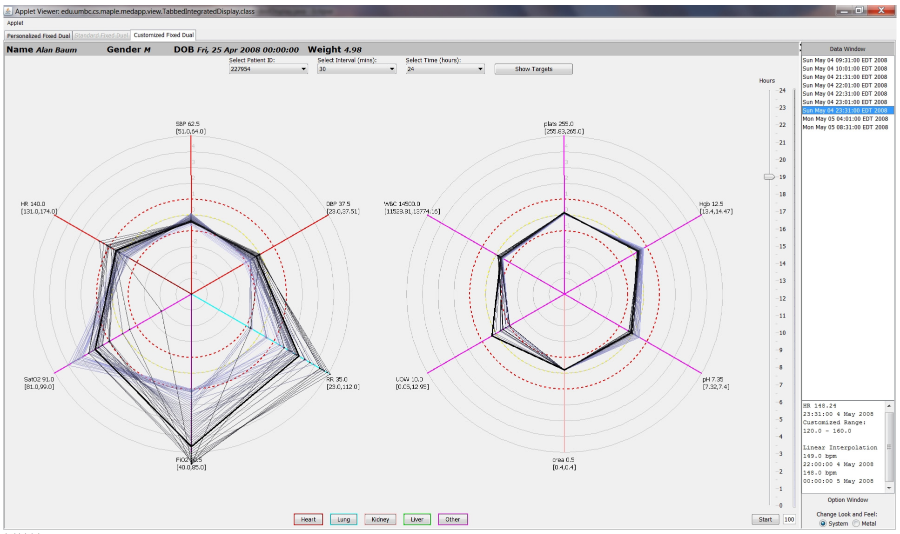

Fig. 4.11 MTSA visualization [102] features two star plots with six variables in each plot. Each polygon connects the values for a 30-minute interval, and can be explored using animation and the vertical slider on the right. The star plot is scaled to the normal range of the variables, whereupon upper and lower bounds are displayed as dotted red circles. The axes are color-coded by the related organ systems.
Image courtesy of Patricia Ordóñez Rozo.

interactions are temporal zooming, panning, and displaying boxes with variable values in place of the glyphs.

The **MTSA visualization** (Figure 4.11) by Ordóñez et al. [100, 101, 102] shows 12 numerical variables of one patient in two star plots [64]. Variable values are aggregated over a time interval (e.g., 30 minutes) and each interval is displayed as one polygon per star plot. It encodes development over time either by animation or by overplotting the star plot with increasing color intensity. The MTSA visualization has two views with different scales of star plot axes: the personalized view normalizes the values of each axis by their arithmetic mean and standard deviation, whereas the customized view scales the values by the normal range of the variables, which the user can customize in the system. Complementing each other, the personalized view emphasizes trends and parameters in a patient, which the user can judge in the customized view for their clinical significance.

In general, visual representations of patient state can use animation to visualize the temporal information of the EHR. For example, Workman et al. [156] use a glyph composed of more than 150 icons, which are positioned in relative anatomical location. Sundvall et al. [137] apply a zoomable map software, Google Earth, to position medical notes inside an illustration of the human body. However, there are trade-offs that need to be considered, when animation is used to encode developments over time. For instance, the comparison between two nonadjacent snapshots is constrained by limited capacity of short-term memory. More animation examples are described below with Gravi++ and TimeRider and further discussion can be found elsewhere [122, 124].

4.2 Visualization of Collections of Patient Records

We now focus on systems that visualize data from multiple patient records at the same time. These systems support quality control and clinical research tasks. On the one hand, they tend to present fewer details about the patient than the systems described above, but on the other hand, they need to provide query methods to find relevant patients. They also need features to aggregate patients to groups, to detect clusters, and to recognize outliers.

4.2.1 Event Sequences

Lifelines2 [147, 148] and Similan [155] are interactive visualization systems designed to search and explore event sequences in multiple records of temporal categorical data. In both Lifelines2 and Similan, each patient record is stacked vertically on a shared horizontal time line, and has its own list of categorical variables. The variable categories are color-coded, and all instances of the same variable category are on the same horizontal line, represented by icons on a zoomable time line.

The distinguishing interaction technique of **Lifelines2** is alignment (Figure 4.12(a)). Users can align all records by a specific event type (for example, heart attacks). Every record's first heart attack event will then be aligned vertically. Records that do not have heart attacks are filtered out. When the records are aligned, the time line switches from an absolute calendar scale—showing the actual date of the events—to

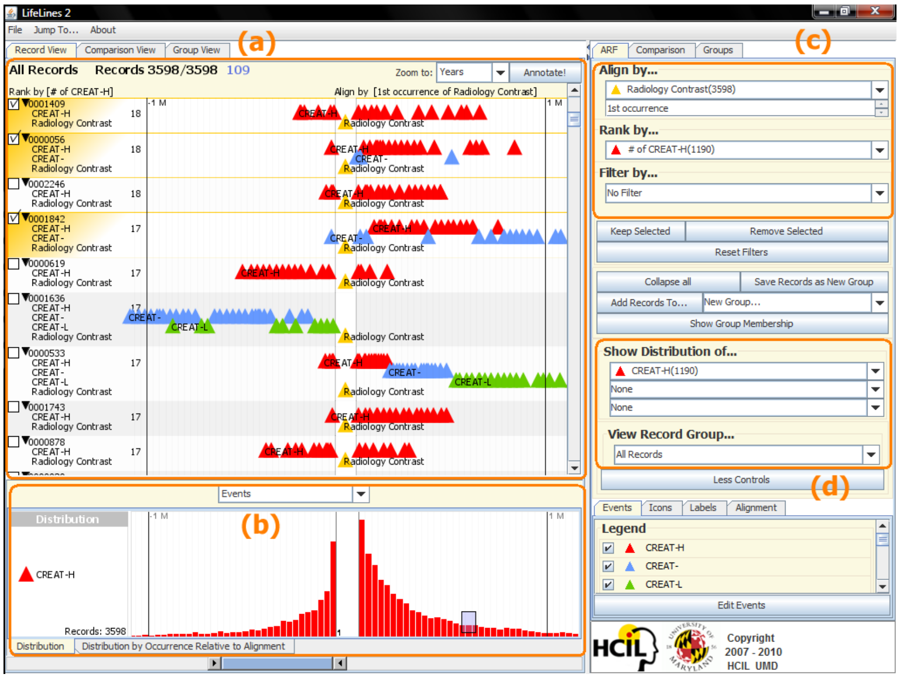

Fig. 4.12 Lifelines2 [148]: visualization of records temporal categorical data. (a) Shows each individual record, and (b) Shows the aggregation of events over time. (c) Contains controls to the basic operators: align, rank, and filter. (d) Shows controls for event distribution and grouping. *Image by the authors.*

a relative scale showing the amount of time before or after the date of the event used for alignment (i.e., one, two, or three days before/after). This allows users to spot trends in the timing of other events relative to the alignment event, in a group of patients. Users can also align by the *n*th occurrence of the event (e.g., the second heart attack). Or they can align by all occurrences. In that case, the display of a record is duplicated for each occurrence of the event and each duplicate is aligned by one occurrence.

Lifelines2 also includes other more traditional operators such as rank and filter (Figure 4.12(c)). Records can be ranked by the number of occurrences of a particular event (e.g., the number of heart attacks or the number of abnormally high creatinine tests). Users can interactively filter the records to find those that exhibit certain temporal patterns of event. They can search for particular sequences of events—including both the presence and the absence of events (e.g., A followed by B, but with no C in between), which is useful in many medical scenarios.

Temporal summaries (Figure 4.12(b)) are histograms showing the temporal distribution of selected event types. Temporal constraints can then be interactively specified on temporal summaries. For example, users can align by the first heart attack, and then draw a box on the temporal summary of the high blood pressure events to select all the records that include high blood pressure events in the six months previous to the first heart attack. Finally, after each filter operation, users can save the results as a new group of records. Subsequent filtering can be applied to iteratively explore different subgroups of patients. In comparison mode, users can compare multiple groups by their temporal summaries. Based on eight medical case studies, Wang et al. developed a generalized process model for Visual Analytics of multiple EHRs [149].

Similan [155] (Figure 4.13) uses the same layout of the records as Lifelines2 and also has alignment capabilities. But it uses a search strategy based on its innovative similarity measure (M&M measure), and ranks records by their similarity to either a selected reference record or to a sequence specified by dragging icons on a timeline. The M&M measure takes addition, removal, transposition of events, and temporal differences of matching into account when computing the similarity of two temporal sequences, and users can adjust the weight of each distance measure component. In a traditional search, records that do not fit the search criteria are removed, so users do not see those that almost fit the specified pattern. In Similan, users can see such results, and can better refine their searches. Combined with alignment, search based on relative time can be performed. Other interactions allow users to specify a time range of interest (absolute or relative).

Both Lifelines2 and Similan concentrate on categorical data, but numerical data can also be visualized if it is first preprocessed (e.g., high/normal/low blood pressure). They do not make any distinction on whether the data is medical test, diagnosis, or treatment. Further those two prototypes focus on point events, as opposed to intervals.

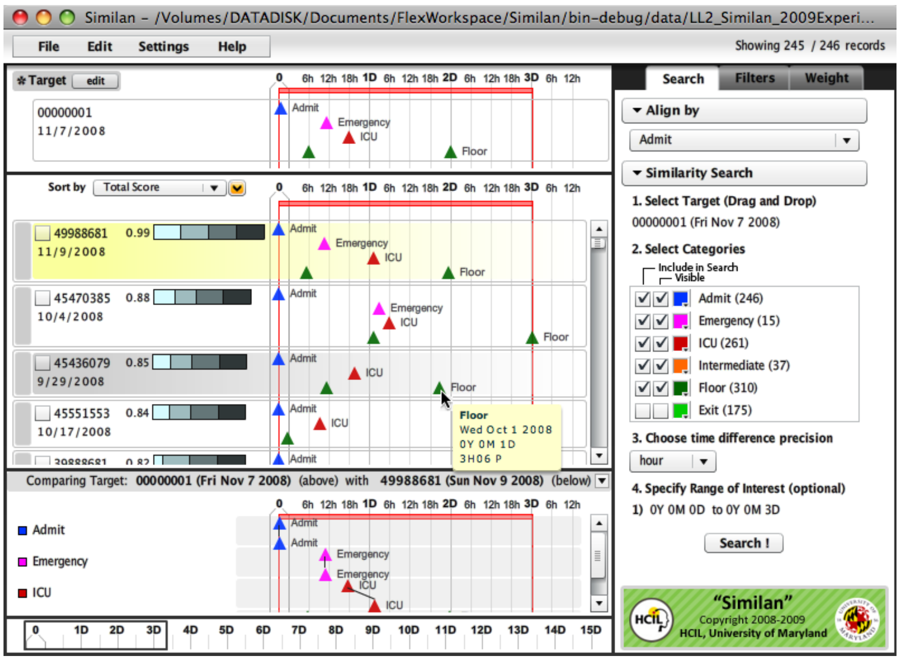

Fig. 4.13 Similan [155]: Similarity search by target (left side top panel). All records are ranked by the similarity match (left side middle panel). Matching detail is shown below. Control panel is to the right. *Image by the authors.*

They would require linking to another visualization to show the details of individual records. Although the number of variables is not bounded, the color-coding scheme and limited short-term memory favor using a small number of variables. Lifelines2 and Similan have been used in case studies with up to 3,958 patients [148]. A user study comparing simplified version of these two systems showed the relative strengths and weaknesses of exact and similarity search [154].

LifeFlow [153], **Outflow** [152], and **VisCareTrails** [88] provide overview visualizations of event sequences extracted from numerous patient records, while representing time spans between events. LifeFlow has been demonstrated with 7,041 patient records in emergency care and 203,214 traffic incidents, OutFlow with 6,328 patient records in a cardiology study, and VisCareTrails with 631 patient records in a cancer study. Furthermore, LifeFlow in combination with the CNTRO

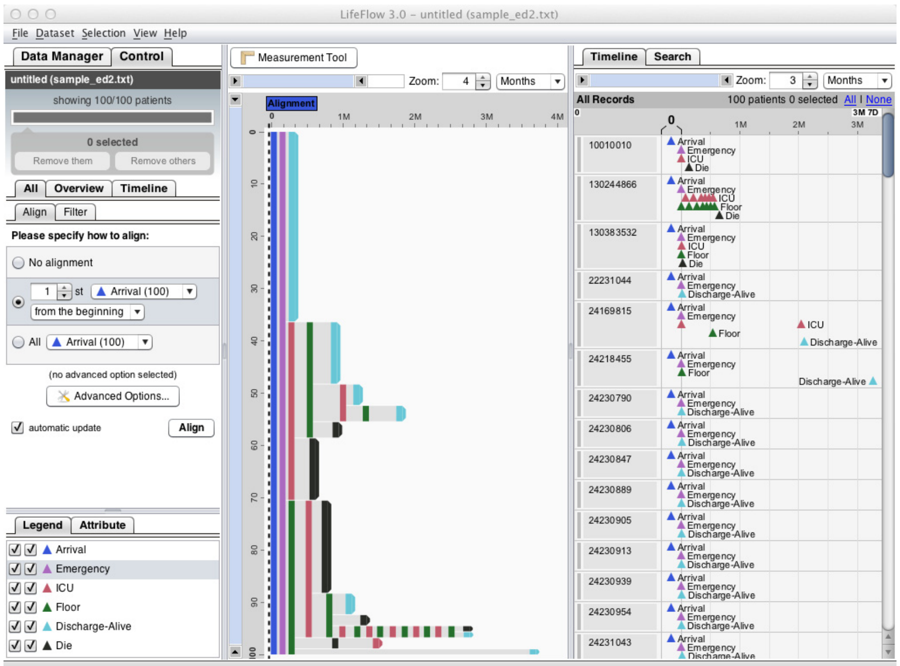

Fig. 4.14 LifeFlow [153]: Overview visualization of event sequences. Here, the center panel shows 100 patients in overview and the right panel displays the first 15 patients in detail. On the left there are controls for alignment, filtering, and selecting event types. In the figure all EHRs are aligned by "arrival", the blue event type. *Image by the authors.*

system exemplifies how data extracted from free text in the EHR can be visually explored [139]. LifeFlow can be combined with Lifelines2 to provide views for both overview and detail tasks (Figure 4.14). LifeFlow and VisCareTrails can be distinguished by the encoding of event types. LifeFlow uses color for a compact layout, while VisCareTrails prints text labels and can thus support more event types. OutFlow uses a graph-based visual presentation and shows the eventual outcome along the event sequences (Figure 4.15).

4.2.2 Expressive Temporal Queries

Fails et al. introduce **PatternFinder** as a search interface for multiple patient records (Figure 4.16) [53]. The novelty of PatternFinder relies on its query specification for temporally ordered events with value

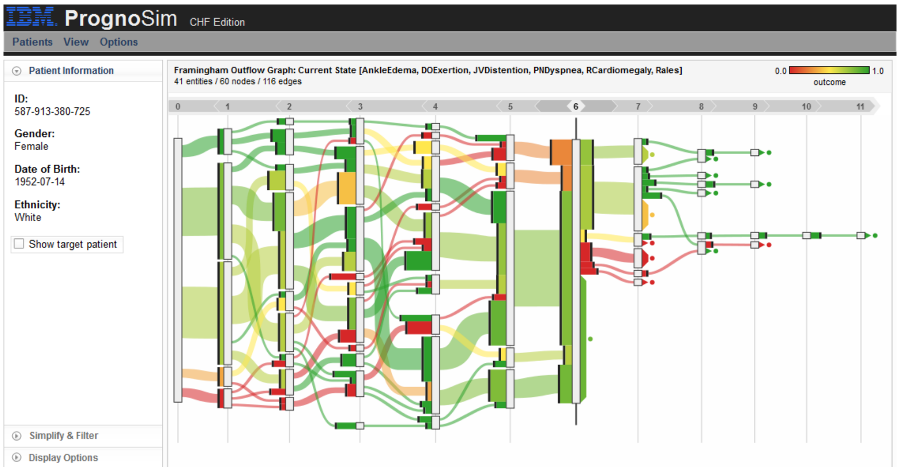

Fig. 4.15 Outflow [152]: Aggregated view of event sequences in a cohort of 41 patients. Colored rectangles represent subgroups of patients split by what event occurred next, or before for layers left of layer 6, which is used for the alignment. The height of each rectangle stands for the number of patients, the width encodes the average time between these two events, and color represents the final outcome averaged on the patient subgroup.
Image by the authors.

and time span constraints. Searching for existence of events, temporally ordered events (e.g., heart attack followed by stroke), temporally ordered values (e.g., 15 or below WBC followed by 16 or higher), and temporal value trends (e.g., monotonically increasing) are all possible. Users can define how far apart each event should be by setting a range of allowable time spans between events. All queries are specified through form-based interfaces.

The underlying data model consists of a higher-level set of categorical events that contain values, which can be numerical or categorical, for example a “White Blood Cell (WBC) count” with a value of 10. The higher-level set of categorical events is organized in a hierarchical vocabulary, for example, “medical tests” may contain the child “blood test”, which in turn contains “blood sugar test” or “WBC test”. Users can choose to use any level in the hierarchy to perform relevant queries.

Finally, PatternFinder introduces a “ball-and-chain” visualization for the search result set. When matches are found, each result record

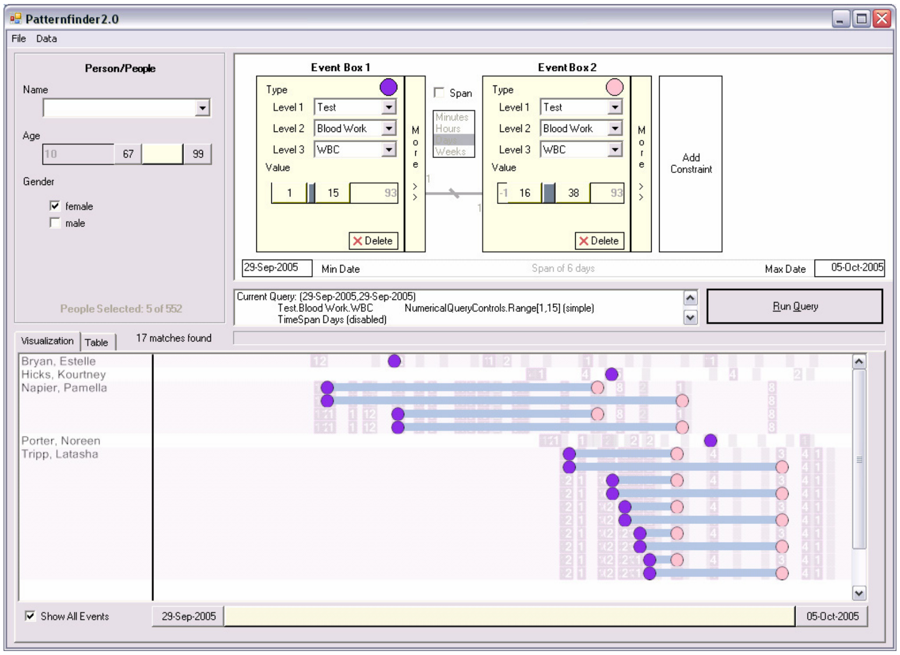

Fig. 4.16 PatternFinder [53]: Top half of the screen shot shows a form-based interface for querying patient records. The results of a query are shown as a ball-and-chain visualization at the bottom. *Image by the authors.*

contains events relevant to the query, which are plotted on a shared time line. Matched events are shown as circles, color-coded the same way they are in the specification interface. Events that match the specified event pattern are linked using horizontal lines. Users can zoom into the time line and highlight specific instances, but otherwise cannot interact with the visualization.

A follow-up design of PatternFinder (Figure 4.17) [109] was prototyped in Amalga, a commercial EHR system from Microsoft. It preserves the ability to specify rich temporal queries with time spans between query elements. Users can also query for relative change of numerical values, for example, 50% drop of platelet count in two days. The interface is tuned for specifying the before and after events for a sentinel event, and is simplified to handle patterns of up to three events (sentinel or reference event, baseline, and follow-on). The result visualization shows all matching patterns automatically aligned by the

Fig. 4.17 PatternFinder in Amalga [109]: The form-based query interface is enhanced by allowing the user to specify an alignment at the time querying. The results are shown with color-encoding and alignment on the specified "sentinel event" (purple markings).
Image by the authors.

sentinel event. Each event is represented by a color-coded tick mark on the time line.

Both PatternFinder versions have more expressiveness in specifying temporal constraints than the sequence search in Lifelines2. However, as the original PatternFinder authors found out, the large number of choices in a form to allow the complex constraints can overwhelm casual users. In the original PatternFinder, the matching events are shown as ball-and-chain, but the other nonmatching events are also shown as gray boxes in the result, giving users context. This context is not preserved in PatternFinder in Amalga: only events that match are shown. While this simplifies the display, it also reduces user's ability to review the results.

The **VISITORS** system by Klimov et al. combines a clinical knowledge base with visualization to enable users to explore multiple clinical records (Figure 4.18) [80, 81]. The system relies on domain ontologies to define clinically meaningful higher abstractions given raw, temporal data. This approach builds on the authors' previous systems

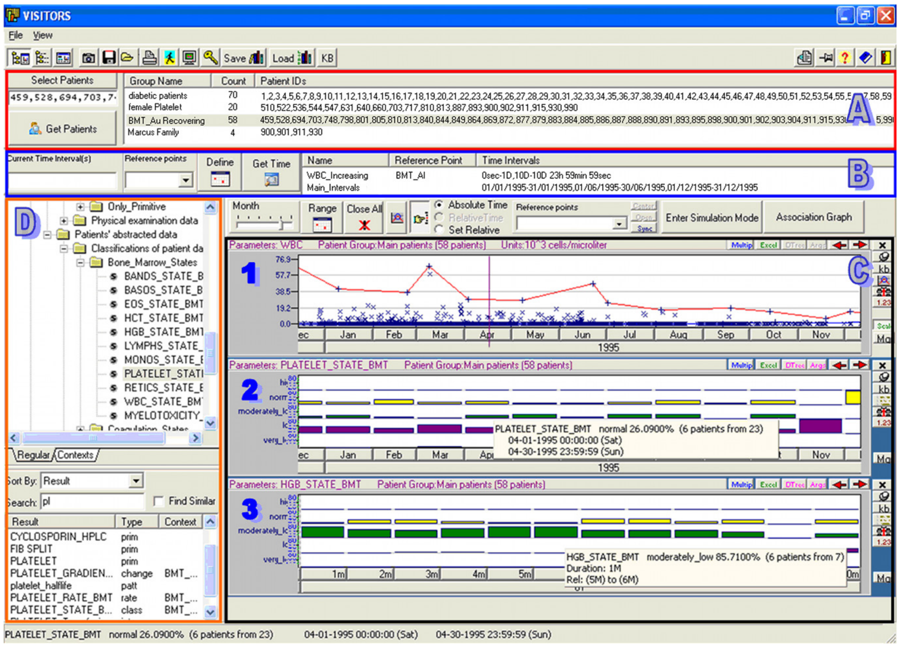

Fig. 4.18 VISITORS interface [81]: (A) shows groups (cohorts) of patients. (B) shows a list of time intervals that are of interest. (C) shows raw data and temporal abstraction of the current groups of patients over the current time interval. Panel 1 shows the white blood cell raw counts for 58 patients, while panels 2 and 3 show the states of platelet and hemoglobin in higher abstraction, respectively. Abstractions are encoded in medical ontologies listed in (D). *Image courtesy of Denis Klimov.*

KNAVE [127], KNAVE II [91, 128]. A similar approach has been proposed in the PROTEMPA system [112]. However, VISITORS is the first system that focuses on groups of patients.

The idea is that while low-level readings of a medical test are important, physicians or clinicians often have to make decisions based on higher abstractions. Given the appropriate domain ontology, these systems take raw numerical data such as white blood cell (WBC) counts over time to automatically derive temporal abstractions such as durations in which patients experienced high, normal, low WBC levels. These temporal abstractions can then be used (sometimes in conjunction with raw data) to define higher abstractions such as Myelotoxicity levels. The VISITORS system builds on top of this and offers expressive ways for users to query for patients.

To show how the derivations are made, the users can see the numerical and categorical raw data as well as all the abstractions that have been derived. These systems typically show raw numerical data as line charts, and categorical data as tick marks (or a bar if it has a duration), both along a horizontal zoomable time line. In VISITORS, these abstractions are shown as aggregations of the values of a group of patients (58 in the example described in the paper).

Differing from other temporal abstraction systems, VISITORS offers an expressive query language to allow users to search for both raw data and abstracted data in groups of patients. Users can specify to search for patients that fulfill one or more of the following: (1) non-temporal constraints, for example, gender and race; and (2) time and value constraints, for example, hemoglobin value less than 10 g/dl and "very-low" platelet value within 3 days after a bone marrow transplant (Figure 4.19). In the time and value constraint specification, a user can specify a range instead of a value. A user can specify combinations of multiple variables, including both raw data and derived temporal abstractions. Finally, the expression can include relative time specifications like "within 3 days after a bone marrow transplant". VISITORS also handles proportion constraints with respect to time (find patients whose WBC levels are at least "high" for more than 70% of the tests in January), and statistical constraints relative to a population (patients whose mean blood glucose is 5 g/dl more than the populations mean). These expressions can also be used to find time intervals instead of patients. The language is simplified by excluding the Boolean operator NOT and nested Boolean expressions. Queries are specified using form-based interfaces so users do not need to learn the syntax of the proposed language. The authors report good accuracy and good usability with the interface and its underlying languages in a usability study among ten persons.

EventFlow [70, 95] builds on the overview and search capabilities of LifeLines2 and Similan while adding interval queries and expanding the specification of the absence of events. Moving to interval data is non-trivial as it requires changes in all aspects of the system but opens the door to new significant clinical research questions such as characterizing

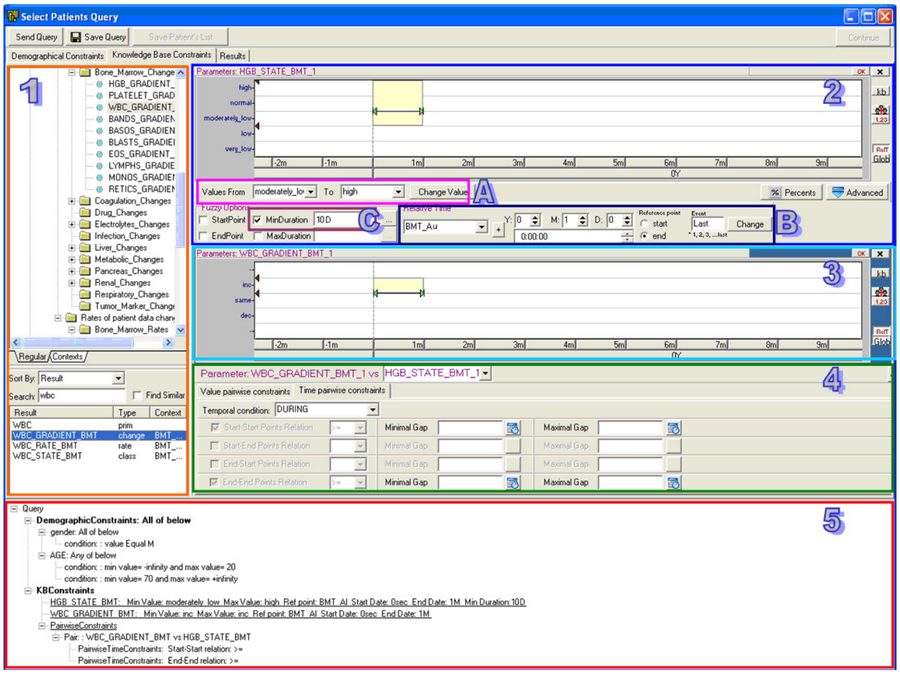

Fig. 4.19 VISITORS's interface for selecting patients [80]: Users can specify temporal and value constraints for two variables through a graphical widget in (2) and (3). Alternatively they can fill out form components (A)–(C). (4) allows users to specify the pairwise temporal relationship between matches for the two variables. (5) shows the generated query and (1) shows the medical ontology. *Image courtesy of Denis Klimov.*

the nature and continuity of medication use. Research conducted with this prototype revealed that the main difficulty for users with questions involving temporal event sequences is not necessarily understanding the underlying complexity, but articulating them into meaningful queries. To address this problem a basic menu-based search interface provides quick access to the most fundamental temporal relationships and serves as an introduction to the advanced graphical-based search interface. In the basic search the subsequence module gives access to the "before" and "after" relationships (e.g., "Drug A" followed by "No Stroke" followed by "Drug B"), while the Overlap module gives access to "during" relationships (e.g., "No stroke" while taking "Drug A" and "Drug B"). The advanced search interface allows users to specify these relationships in tandem, as well as access more complex temporal features such as absolute time constraints and the full range of absence scenarios. The advanced search interface revolves around a visual query language that is used to draw complex sequence of event relationships, including intervals and absence of events (Figure 4.20).

Other systems allow expressive temporal queries over EHR databases through visual metaphors: **VizPattern** [74] uses a comic strip metaphor, for example, two icons in the same panel denote simultaneity. Combi and coauthors present a **Paint Strips** metaphor [44, 50] that includes different metaphors for temporal granularities.

4.2.3 Patient Records on a Common Time Axis

Brodbeck et al. [36] introduced **Caregiver** to support therapeutic decision-making (Figure 4.21). Caregiver has three main visualizations, all on the same horizontal time line. On the top panel, it visualizes different cohorts of patients (grouped by, for example, different intervention methods). Each cohort is represented by a solid rectangle on the timeline (x-axis), where the length indicates the duration of data for that group of patients and the height indicates how many patients are in that cohort. On the bottom panel, Caregiver shows an overview of all patients. Selecting a cohort highlights patients in that cohort, and vice versa. Each individual patient occupies a horizontal strip of space, allowing the display of a single chosen variable in a bar chart. Finally, users can select a particular patient to expand and show more variables in line plots or larger bar charts. In this respect, ordinal and numerical variables are most appropriate for Caregiver. Different intervention methods can be implicitly represented in the patient cohort group, but interaction with them is limited to selection and sorting. Furthermore, it provides a focus and context technique for the time axis, so that a time period in focus is zoomed, while the remaining time periods are shown compressed. The system has been demonstrated to work with 100 patients, over 60,000 observations, and with more than ten variables.

Caregiver is primarily a search interface for queries by patient status but not for developments over time. Users perform searches in one of

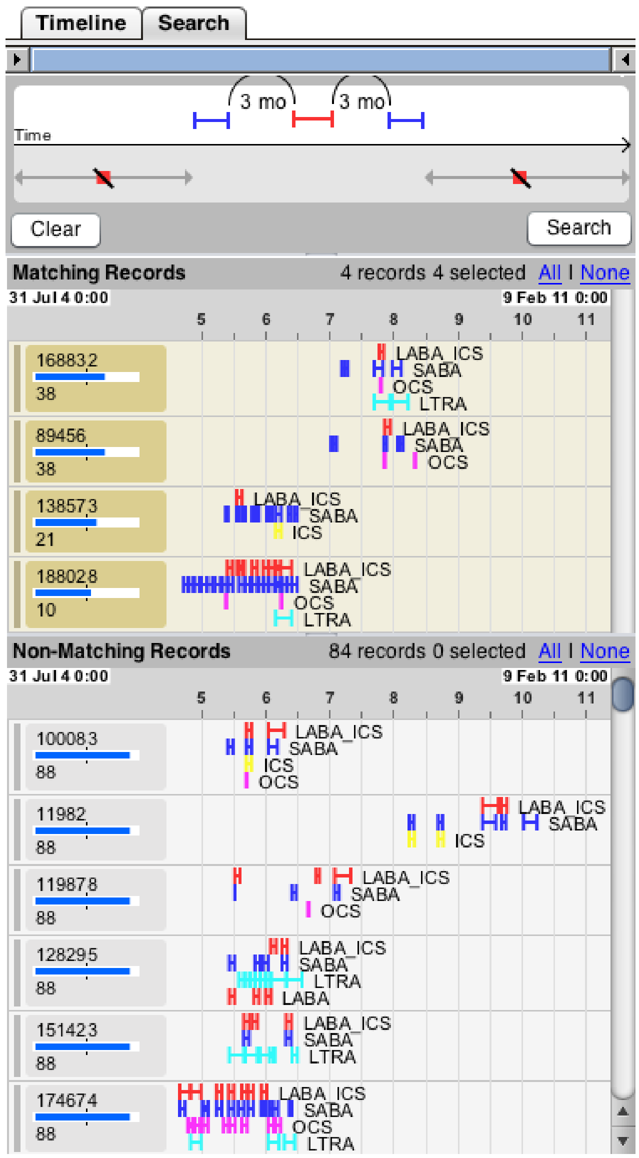

Fig. 4.20 EventFlow [95]: Users can interactively compose complex temporal queries with intervals. Here a user was looking for patients who received a weaker asthma medication ("SABA", in blue) within 3 months of a stronger asthma medication ("LABA_ICS", in red). They also wanted to ensure that this sequence was neither preceded nor followed by the strong medication. Search results appear highlighted directly below the visual query, while nonmatching records are still visible in the bottom to help check that the query was executed as expected. *Image by the authors.*

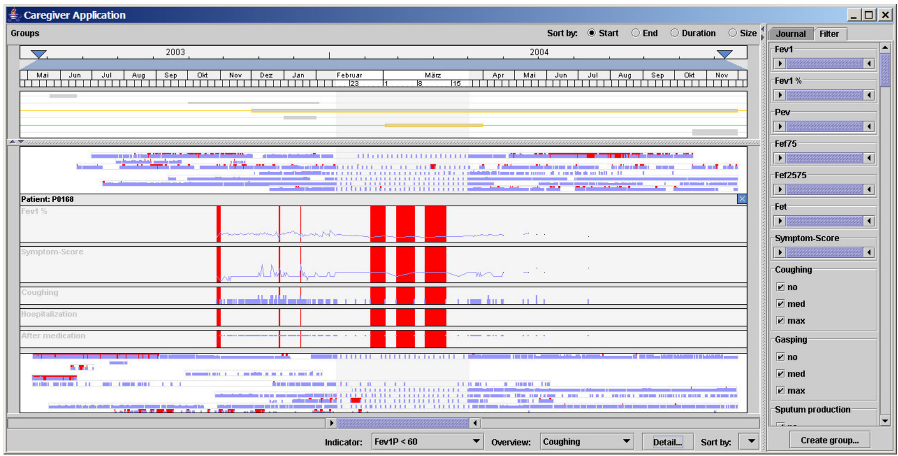

Fig. 4.21 Caregiver [36] displays both numerical and categorical data on a horizontal timeline with a variety of search and highlight options. Temporal values that satisfy a search criterion (here Indicator $Fev1P < 60$) are highlighted with a red background. The time period from February 20 to March 18 is in the zoomed focus area.
Image courtesy of Dominique Brodbeck.

two ways. First, they can use control widgets to dynamically query for patients who match certain values in the variables. Those who do not have these values are grayed out or excluded from the display altogether. Second, an "indicator" can be defined to highlight temporal regions that exhibit the defined characteristics (the red regions in Figure 4.21). An indicator can be a combination of values of any number of variables. Based on these search results, users can create new cohorts. These search options allow users to find time periods where values of variables conjunctively occur, but not allow search of temporal patterns such as high coughing symptoms followed by high temperatures.

Chittaro et al. [45] presented **Interactive Parallel Bar Charts (IPBC)**, a 3D visualization of numerical data from multiple hemodialysis sessions (Figure 4.22). Each numerical time series is shown as a 3D bar chart. One of the horizontal axes represents time, and bar height represents the values of the time series. Multiple bar charts are lined up on the third axis, allowing users a view of all time series at once. The main usage of IPBC is to visualize time series of a single variable (e.g.,

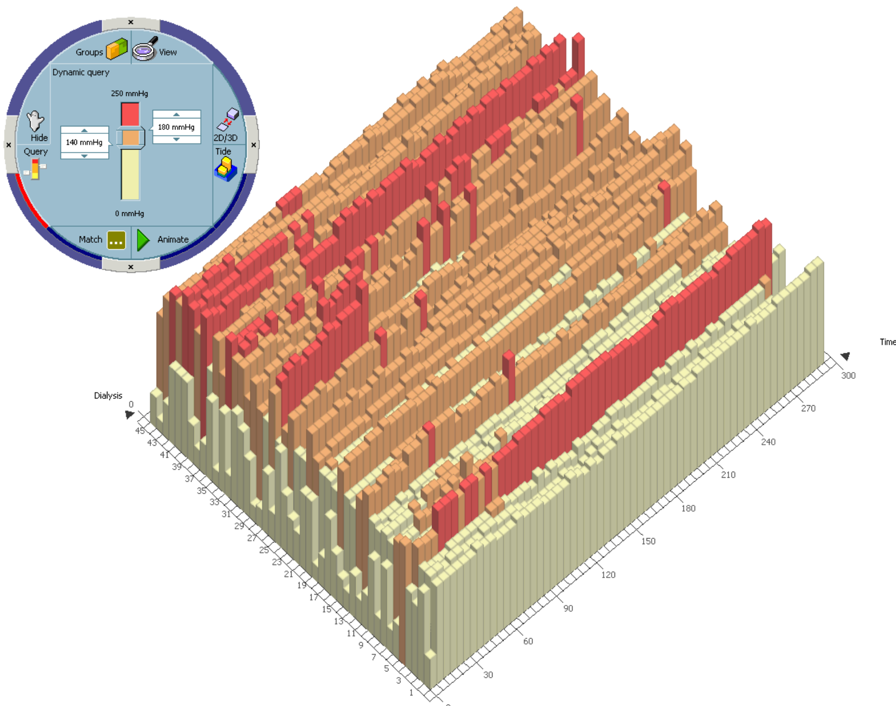

Fig. 4.22 Interactive Parallel Bar Charts (IPBC) [45]: 3D bar chart for evaluation of dialysis sessions. Time is running from the bottom center to the upper right and individual dialysis sessions are aligned one after another to the upper left. Colors are used to signify different qualitative value regions as specified in the round user interface in the upper left corner. *Image courtesy of Carlo Combi, adapted by the authors.*

systolic blood pressure) across multiple sessions of the same patient, where the beginning of each time series represents the beginning of each hemodialysis session. However, IPBC can just as well be applied to multiple patients instead of multiple sessions. Even though only one variable is shown over time across all sessions at a time in IPBC, multiple IPBC displays can be coordinated to allow users to examine the relationships among different variables. Medical test values such as blood pressure and blood flow are used to demonstrate the system. In the paper, 19,000 points of blood pressure data over 79 individual hemodialysis sessions for a single patient are used as an example.

The bars in the IPBC can be color-coded by four value regions: high, normal, low, and out-of-range. These value regions are dynamically controlled by users and allow values of interest to be visually distinguishable. IPBC allows users to define a linear threshold function with respect to time and highlight values above the threshold. Time series data can be aggregated to show mean values to reduce visual noise. For example, values can be averaged in 10-minute bins and shown as an average. IPBC offers traditional 3D controls to zoom, rotate left/right, and rotate up/down. To overcome problems of 3D occlusion, users can flatten the entire (or part of the) display, showing the values of each bar as a color-coded cell in a matrix. IPBC allows additional space to be inserted between time series to sacrifice information density in order to reduce occlusion.

Users can use query-by-example to find similar patterns of values across time series. By selecting a region of interest on a time series, IPBC can highlight regions of all time series that are similar to the target region within some specified tolerance. If the task requires understanding relationships among multiple variables, users can create multiple IPBC views, each focusing on a single variable, and explore them using coordinated highlighting. This approach uses a lot of screen space, so more than three variables at a time may be difficult. Alternatively, the authors offer a pairing of IPBC to a parallel coordinates plot [71], which visualizes multiple variables at the same time. Using the parallel coordinates plot, users can comfortably study relationships among six or more variables at once while also focusing on the important one in IPBC. In the parallel coordinates plot, however, the time dimension is only visible as an animation.

CareCruiser [62] supports the analysis of the evolution of patients' condition in response to treatments and emphasizes the effects of treatments on the patients. The center part of its user interface (Figure 4.23) consists of a set of panels containing plots of numerical variables. Treatment plans are represented by colored rectangles placed below the plots, and diamond shape icons indicate the timing of individual actions taken during the treatment. A distinction is made between planned actions and unplanned actions, which are displayed below the others. To compare the effects of treatment, the plots are aligned by the start of a treatment plan or by a specific action, so that different patients can

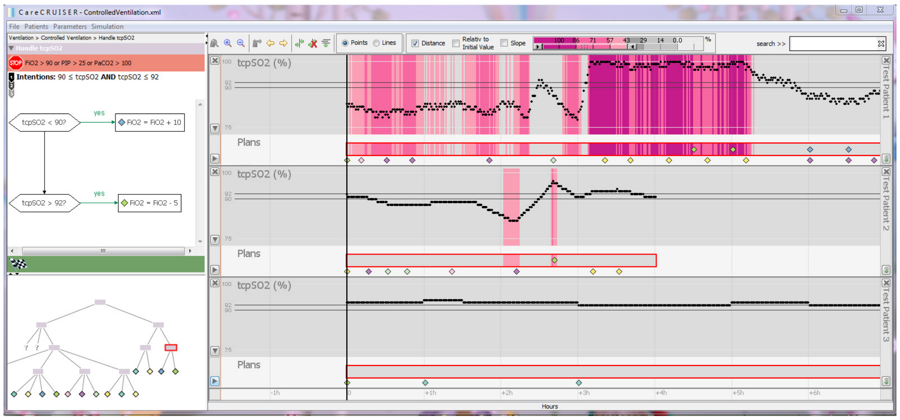

Fig. 4.23 CareCruiser [62]: Visualization of medical actions along with their effect on patients' condition. Here, the point plots show oxygen saturation of three patients and the diamonds below each plot represent planned and unplanned actions. The pink background emphasizes distance from the intended value range. Patient panels are aligned by the beginning of the selected treatment plan, which is shown by a vertical black line.
Image courtesy of Theresia Gschwandtner, used under CC-BY-ND license.

be compared on a relative timescale. Alternatively, users can align different instances of the same clinical action in the data of one patient repeated in several panels. Furthermore, each panel may contain point plots for more than one variable.

The background color of the panel emphasizes the effects of clinical actions and the progress of treatment by three different abstractions: distance to goal value (cp. Figure 4.23), slope (rising/falling of the value), and progress from initial value to the intended value.

VisPap [135] is a multiple coordinated views system for heterogeneous data in cohort studies, for example in neuropsychiatry. Its datasets includes both features extracted from medical images and laboratory data. It visualizes the evolution of one variable over time in a line plot, which allows alignment by calendar date, patient age, and first occurrence of the variable. Scatter plots and parallel coordinate plots allow users to explore the mutual variations of different variables. It also provides renderers for volumetric images, so that users can drill down on distinct features found in the plots.

InfoZoom [133] was used to visualize large tables of medical variables from a patient cohort. Aggregated views show the distribution of values in large number of patients on a single screen. Progressive filtering and sorting allow rapid exploration of those distributions and access to details.

4.2.4 Snapshots of Multivariate Patterns in Patient Cohorts

Gravi++ and TimeRider represent patients as marks that are spatially arranged based on their variable values. Thus, patterns such as clusters and outliers become visible in the patient population. Animation and traces represent the time dimension.

Gravi++ [67] uses interactive visual clustering to allow exploratory analysis of multiple categorical variables across multiple patients over time (Figure 4.24). Each variable is mapped to an icon, and each patient is also mapped to an icon (of a different type). The placement of the patient icons is dependent on a spring-based layout [33]. That is, the higher the value a variable is for a patient, the closer that patient's icon is to the variable icon. The idea is that patients who share similar values in the variables will then be placed in similar locations on screen, allowing users to visually detect clusters with ease. Since it is hard to read variable values from the spring-based layout, Gravi++ displays each patient's parameter values as circles around variable icons—a feature called attraction field.

Because the value space determines the placement of the patients, this system is ideally suited for ordinal variables. The visualization of these values is the most important aspect of Gravi++, but patient attributes can be also encoded by patient icons, for example, size can encode age or body mass index, and color can encode gender or therapeutic outcome. There is no known hard limit on the number of variables the system can handle, but too many variables can make it hard to find a meaningful configuration. The paper demonstrates Gravi++'s effectiveness with four medical variables and one patient attribute for nine patients over six time steps.

Users interact with the system by changing the encoding of the variable icons. They can select which variables to show or drag their

Fig. 4.24 Gravi++ [67]: spring-based visualization for evaluation of psychotherapeutic measures. The six square icons on the perimeter represent variables. The placement of the circular patient icons in the middle is based on each patient's values in the six variables. Furthermore, green and red circles around the variable icons show variable values of the three patients. On the left, users can choose patients and variables from a list.
Image by the authors.

icons into different locations to examine their influence on the patient population. Users can also modify the spring force of the variable icons. Users can superimpose a star plot [64] on top of the Gravi++ display to get a one-glance view of all patients' values at the same time point. When a patient is selected, the circles around variable icons showing the patient's values are highlighted. Finally, users can remove or add patients to the visualization.

Gravi++ uses animation or traces to represent time. Users can view the patient icon positions for a single time point, play or step through all time points. Alternatively, users can activate traces, which show the patient icons of all time points connected to a polygonal line. Following the position of the patient variable icons through the animation allows users to explore, for example, factors that affect the forming of clusters. The authors demonstrate Gravi++'s success in allowing domain experts to find predictors of success in psychotherapy for anorexia nervosa patients.

TimeRider [122] employs an extended animated scatter plot to visualize trends in patient cohorts (Figure 4.25). In contrast to Gravi++, it places patient marks along two axes that can encode numerical or categorical variables. Up to three additional variables can be mapped to color, shape, and size of marks. On the one hand, this limits TimeRider to showing bivariate patterns compared

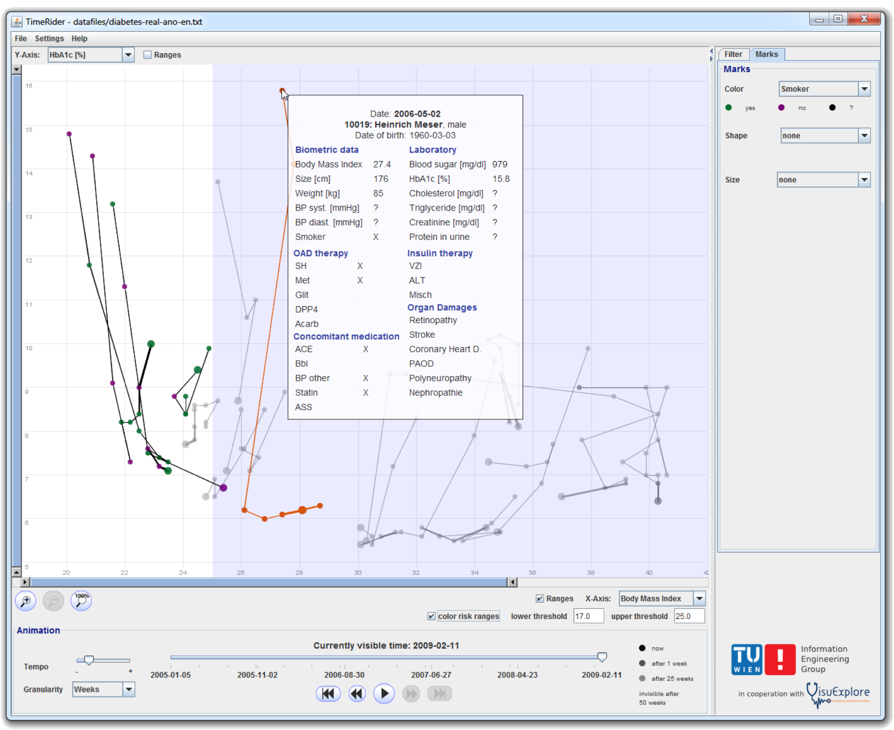

Fig. 4.25 TimeRider [122]: Animated scatter plot with optional traces. Patient details are shown in a semitransparent overlay when hovering over a data point. Color is used to show smoking (green) and nonsmoking (purple).
Image by the authors, used under CC-BY-ND license.

to multivariate patterns in Gravi++, but on the other hand, the positions in TimeRider are immediately meaningful.

TimeRider represents time by either animation or traces. To handle irregular sampling in animation, transparency is used to encode data wear. Synchronization modes offer different views on the patient cohort: calendar date, time since start of therapy, time before end of therapy, and patient age. The areas of the plot background can be filled with color to emphasize variable ranges that are critical for patients. Furthermore, the system supports interactions to change variable mapping, pan and zoom, select or filter patients, and show a tooltip.

The **Dynamic Icons (DICON)** system [61] clusters EHRs that are similar to a target record. Then it visualizes these clusters as composite icons that have parts representing the features of all EHRs in the cluster. It provides different spatial arrangements of the icons (e.g., in a scatter plot) and manual refinement of the clusters.

5

Discussion

The systems in our survey demonstrate a broad spectrum of visualization and interaction techniques to deal with large amount of complex EHR data. First, we analyze the strengths and weaknesses of the 14 systems selected for detailed review, using the criteria introduced in the *Methods* section. Next, the results of evaluation studies performed in a medical context are summarized. Finally, limitations and future directions are discussed.

5.1 Analysis of Review Criteria

Tables provide an overview of the characteristics of the systems (Table 5.1 for data types, number of variables on the screen, and EHRs of multiple patients; Table 5.2 for medical information types and scenarios; Table 5.3 for support of user intents).

5.1.1 Visualization of Categorical Data

The most common way to visualize categorical data is by placing icons (for point events), and line segments (for events with duration) on a horizontal time line (LifeLines, Midgaard, Lifelines2, Similan, PatternFinder, and Caregiver, VISITORS). Most systems separate

Table 5.1. Summary of review criteria for 14 systems.

		categorical data	*numerical data*	*no. of variables per screen*	*one patient*	*multiple patients*
single EHR	LifeLines	●	○	~ 25	●	
	MIVA	○	●	~ 5	●	
	WBIVS	●	●	10	●	
	Midgaard	●	●	~ 15	●	
	VisuExplore	●	●	~ 10	●	
	VIE-VISU	●	●	15	●	
EHR collection	Lifelines2	●		~ 10		●
	Similan	●		~ 10		●
	PatternFinder	●	○	3		●
	VISITORS	●	●	~ 5	●	●
	Caregiver	●	●	1–6		●
	IPBC	○	●	~ 3	●	●
	Gravi++	●	○	~ 6		●
	TimeRider	○	●	2–5		●

●: Full support, ○: partial support, “ ”: no support.
The number of variables per screen is an estimate based on examples in the original publication.

different categories of events on separate bands. They often use color to encode different types of events, and allow users to customize icons and colors. Some provide shape and size encoding schemes for icons. In LifeLines, semantic zooming is used to merge/separate closely related categories (e.g., drug A and drug B may be merged into a single drug class when users zoom out).

Differing from all other systems, Gravi++ uses positions of patients to indicate the values of categorical variables, and animation to indicate time. This way, Gravi++ makes the clustering of patients by variable values easier to see.

5.1.2 Visualization of Numerical Data

For systems that handle numerical data, using line plots is the most common approach (e.g., Midgaard, WBIVS, Caregiver, VISITORS). Point plots and bar charts are also used, sometimes in conjunction with line plots (e.g., VisuExplore). TimeRider encodes the two axes of the scatter plot to two numerical variables. In addition, the size of graphical elements can be mapped to a numerical variable (e.g., VIE-VISU, Gravi++).

5.1.3 Combined Visualization of Categorical and Numerical Data

Most systems that visualize both categorical and numerical data make use of multiple representation techniques and align them along a shared time axis. For example, WBIVS primarily deals with numerical data using line charts, but uses a matrix view to show categorical variables. CareCruiser renders numerical test results in a point plot and medical actions as diamond symbols.

VIE-VISU is special in several ways. First, it uses small multiples to indicate time progression, and does not require a time axis. Instead, it uses sizes of elements of a glyph to indicate a number of numerical medical values at once. It also uses color-encoding to show the values of ordinal variables. TimeRider maps two numerical variables to the position of the patient marks and up to two categorical variables to color and shape of the marks, while animation is used to encode development across time.

Some systems build categorical abstractions of raw numerical data. VISITORS shows numerical data as a combination of point plots and line charts while categorical abstractions are shown as size and color-coded rectangles using the same time line. Midgaard provides semantic zooming so that as users zoom in, color-coded rectangles change to a bar chart and then into a line chart to show more detail. Similarly, IPBC can collapse its 3D bars into a matrix plot to get an overview. Both the bars and the matrix cells are color-coded to represent the categories “high”, “normal”, and “low”.

PatternFinder, VISITORS, Caregiver, and TimeRider can be used to formulate queries that combine categorical and numerical variables.

5.1.4 Medical Information Types

In general, the systems we found do not use distinct visualization and interaction techniques for specific types of medical information (e.g., diagnoses, treatments, or drugs). On the other hand, medical test results tend to be numerical while diagnoses and treatments tend to be categorical, and whether the data is numerical or categorical has some impact on the visualization. Most of the systems deal with medical tests (e.g., blood pressure, white blood cell count), but some systems require categorizing the results as normal/abnormal before visualizing the data. While it is essential to help physicians make accurate clinical decisions by providing them with detail access to test results (such as in MIVA, VisuExplore, or VIE-VISU), other systems designed for hospital quality assurance or clinical research may instead need access to higher-level abstractions (e.g., in VISITORS), or require only categorical data such as low, normal, and high as provided in Lifelines2. Table 5.2 highlights the specific ways systems are tailored to deal with these different medical information types.

LifeLines groups related items in facets. For example, all medical tests belong to one facet, and all medical treatments belong to another facet. Because facets are collapsible, users can expand only the facets that are important to them and avoid using screen space for irrelevant facets. Likewise, the VIE-VISU glyph is composed of three parts representing three physiological systems. On the other hand, some systems use a categorization schema to control how data are accessed and show only a subset of the data. PatternFinder organizes the data in a schema. Users issue queries using such schemas, in which medical tests, diagnoses, and treatments are separated.

5.1.5 Support for Multiple Variables

The most common approach to deal with multiple medical variables is to place them along the same horizontal time axis (e.g., LifeLines, MIVA, Midgaard, VisuExplore, Lifelines2, Similan, Caregiver, VISITORS). A shared time axis facilitates seeing temporal relationships between variables. Typically, visual representations of different medical variables are placed in separate panels, in order to reduce clutter and to accommodate for distinct scales and value ranges (cp. [111]). A different approach puts each variable on its own time axis (IPBC), and shows them side by side. While this makes the temporal comparison of

Table 5.2. Medical information types and medical scenarios that have been demonstrated on 14 systems.

		Tests	*Diagnoses*	*Treatment*	*Details*
Single EHR	LifeLines	●	●	●	Events and intervals for diverse medical information
	MIVA	●		●	Tests and treatments recorded in intensive care
	WBIVS	●			Pulmonary function and subjective symptoms
	Midgaard	●	●	●	Tests and treatments in intensive care and treatment plans
	VisuExplore	●	●	●	Tests, concomitant diseases, and treatments in chronic disease care
	VIE-VISU	●		●	Circulatory, respiratory, and fluid balance plus ventilation settings
EHR collection	Lifelines2	●	●	●	Test, diagnoses, and treatment events. Numerical test events needs to be first converted to categories.
	Similan	●	●	●	Test, diagnoses, and treatment events. Numerical test events needs to be first converted to categories.
	PatternFinder	●	●	●	Test and treatment events (e.g., creatinine and contrast administration)
	VISITORS	●	●	●	Mostly test and treatments data. Both numerical and ordinal are possible. Diagnoses can be implicitly encoded by cohorts.
	Caregiver	●	○	○	Pulmonary function, subjective symptoms, and treatment groups
	IPBC	●		●	Tests and treatments recorded during dialysis sessions
	Gravi++	●	○	○	Questions and indicators in cognitive behavior therapy
	TimeRider	●	○	○	Tests, concomitant diseases, and treatments in cohorts of long-term diabetes patients

●: full support, ○: partial support, " ": no support.

multiple variables more difficult, coordinated exploration via brushing and linking (e.g., in IPBC) mitigates some of the difficulty. IPBC also integrates parallel coordinates to deal with multiple variables. Caregiver allocates most of the display space on one variable, but allows multivariate queries.

Systems like VIE-VISU and Gravi++ show snapshots of the EHR data instead of aligning it to a predominant time axis. VIE-VISU's glyphs encode a number of variables by size and color of the glyph elements. These elements in the glyphs represent an aggregation of values over one-hour span, and the interaction of multiple medical variables over time may be difficult to see. Gravi++ and TimeRider use layout to cluster patients based on two or multiple medical variables. Animation and traces are used to show the history of the variable values over time. While Gravi++ supports an indefinite number of variables, the clustering works best for patients who have similar values in multiple variables.

5.1.6 Support for Multiple Patient Records

For systems that deal with collections of patient records, the most common practice is put all records on parallel lines sharing the same time axis (e.g., Lifelines2, Similan, PatternFinder, Caregiver, IPBC, VISITORS). IPBC uses a 3D coordinate system instead of 2D. VISITORS overlays all patients' data points on the same coordinate space. Lifelines2 additionally provides temporal summaries to show temporal distributions.

Differing from single patient systems, these systems emphasize querying, sorting, aggregating, and clustering of the patients. However, systems vary in the expressiveness of their query languages and approaches to aggregation.

Some systems use form-based user interfaces to search and filter. Caregiver uses both dynamic query and a threshold indicator to perform filtering but without temporal constraints. VISITORS and PatternFinder provide expressive but fairly complex user interfaces to formulate queries. In Lifelines2 users can search using series of align, rank, filter, or group operations or specify temporal constraints using direct manipulation on aggregated views. It allows temporal sequence search, including the absence of events. Finally, IPBC and Similan use query-by-example to allow users to select patients with similar patterns.

A second distinction is whether or not systems support operations on subgroups of patients. Lifelines2 and Caregiver allow the creation of new groups based on the results of filtering, and Lifelines2 allows the groups to be compared with parallel temporal summaries. Gravi++

and TimeRider can color-code patient marks by patient attributes, and create visually distinctive patient groups, but provide no additional ways to manipulate the groups.

A third distinction is whether the systems support aggregation and clustering of groups of patients. Lifelines2 aggregates categorical events over multiple patients in histograms. Similan sorts patients by similarity to a target temporal pattern. TimeRider and Gravi++ position patients, so the ones sharing similar values are clustered together. LifeFlow, Outflow, and VisCareTrails aggregate patients with similar event sequences.

5.1.7 Support for User Intents

While all applications are designed for interactive exploration, the foci and depths in respect to user intents differ as presented in Table 5.3.

Looking at explorability (i.e., the *Explore* user intent) reveals system differences since it plays a leading part in most of the systems, especially in single patient systems. As all systems use temporal aspects of the medical data as a frame of references, all but two systems support navigation in time. Manual pan and zoom on the X-axis are typical interaction techniques for this (e.g., LifeLines). Alternatively, Gravi++, TimeRider, and MTSA use animation and provide play, step forward, and step backward interaction techniques.

For multiple patient systems the filtering options (i.e., *Filter* intent) to show or highlight something conditionally is of higher importance, because these systems need to work with large EHR databases. Caregiver demonstrates a typical dynamic query interface for filtering by patient status. Lifelines2 can query EHRs by event sequences and PatternFinder allows the formulation of more complex queries with relative change over time and temporal constraints. In contrast to these form-based query interfaces, IPBC supports the filter intent visually, for example through query by example. LifeLines is the only single patient system catering to the filter intent. Its search function highlights all items with an implicit relationship with the search term (e.g., “diabetes” would highlight all blood sugar tests and insulin prescriptions).

Table 5.3. User intent support for interactive exploration and querying of EHR.

		Select		*Explore*			*Reconfigure*				*Encode*		*Abstract/ Elaborate*			*Filter*			*Connect*		
		Keep track	*Manage groups*	*Navigate in time*	*Add/remove parameters*	*Add/remove patients*	*Reposition items manually*	*Sort items*	*Adjust axis*	*Other techniques to avoid occlusion*	*Switch representation technique*	*Vary visual encoding*	*Parameter abstraction*	*Temporal data binning*	*Show details of items*	*Patient status*	*Development over time*	*Development with time constraints*	*Patient/group relationship*	*Brush in other representation*	*Brush other parameters*
Single EHR	LifeLines		n.a.	●	○	n.a.						●	●	●	●	●			n.a.		
	MIVA	●	n.a.	●	●	n.a.	●							●	●				n.a.		●
	WBIVS	●	n.a.			n.a.							○	●	●				n.a.		●
	Midgaard		n.a.	●		n.a.	○		●		●		●	●	●				n.a.		
	VisuExplore	●	n.a.	●	●	n.a.	●	○	●		●		○		●				n.a.		
	VIE-VISU		n.a.	●		n.a.					○			●					n.a.		
EHR collection	Lifelines2	●	●	●	●	●		●	●	●	●	●		●	●	○	●	○		●	
	Similan	○		●	●			●	●	●		●			●	○	○	○			
	PatternFinder			○	○					●	○				●	●	●	●			
	VISITORS		●	●	●	○			●				●		●	●	●	●			
	Caregiver		●	●				●	●		●	●			●	●			●		
	IPBC	○			●					●	●		●	●		●	●	○			●
	Gravi++	●		●	●	●	●		●	●	●	●	●		●					●	●
	TimeRider	●		●	●						●	●			●	●					

●: Full support, ○: partial support, “ ”: no support, n.a.: not applicable for single-patient systems.

Abstract/Elaborate, that is, showing less or more details, is the third intent to play a leading part in the reviewed systems. Almost all of them provide a function to show details about items, typically in a tooltip. Parameter abstraction and temporal binning are frequently employed to fit larger data volumes in the display or make them more easily comprehensible. For example, LifeLines merges events into summary events, when needed. Midgaard showcases parameter abstraction by using smooth semantic zooming at interactive speed.

The *Select* intent, for example, marking some items as interesting, is only supported by some systems. A reason might be that selection is often regarded as a way for users to perform additional manipulations on the selected data (cp. [157]). In the absence of such operations, having selection seems superfluous. Nevertheless, keeping track of items for a short term or in groups can be helpful, as we can see with Lifeflines2, VisuExplore, WBIVS, or Gravi++.

The support for the *Reconfigure* intent, that is, showing a different arrangement, largely depends on the chosen visualization method. In general, systems should consider providing interaction techniques that allow end-users (or possibly system administrators) to adapt the visual layout to their needs. With Gravi++, users can reposition variable icons either freely or on a circle. Aligning patients by a selected event helps comparing patient histories and is possible in most systems for multiple EHR (e.g., Lifelines2, Similan, CareCruiser, TimeRider). Likewise, sorting patients by events (e.g., Lifelines2, Similan) or parameter value (e.g., Caregiver) can boost lookup tasks.

Given the wide range of possible visual mappings, it is surprising that the reviewed systems do not more widely support the *Encode* intent, that is, changing the way the data is represented. Notable exceptions are VisuExplore, Gravi++, AnamneVis, Lifelines2, and IPBC. VisuExplore provides different visual representation techniques for numerical data (line plot, bar chart, semantic zoom chart, horizon graph). Gravi++ can switch between three representation techniques (icons, rings, and star plot). AnamneVis has multiple coordinated views for temporal, anatomical, causal, and ontological aspects of the data. Lifelines2 integrates a histogram display and IPBC includes a parallel coordinates plot. Further, Midgaard's semantic zoom smoothly morphs between different representation techniques. Even if the representation technique should remain stable, the users could vary some visual attributes. For example, in Gravi++ users can map variables to the color, size, and label of patient icons.

The *Connect* intent, that is, showing related items, is not widely supported in the systems reviewed. The most common interaction technique is brushing items in other representations or items for other variables at same point in time or the same patient. LifeLines' search function illustrates that EHR data contains many implicit relationships (e.g., tests and medications can be linked to diagnoses). Increases of white blood cell counts could be linked to an infection, or concepts could be related through an ontology, but such relationships were not conveyed in the visualizations we reviewed.

5.2 Empirical Evaluation in Medical Context

The evaluation of visualization methods and systems encompasses studies on their effectiveness, efficiency, and usability. These are important for scientific research and practical application. Empirical evaluation of visualization can be challenging [41], particularly for a complex process like medical decision-making. There are multiple possible evaluation methods, which cater to different research objectives and have trade-offs in terms of precision, generalizability, and realism [41, 84]. Furthermore, it is often difficult to recruit medical professionals as test users. Therefore, some studies were conducted with subjects that do not have a medical background (e.g., Midgaard [27], Gravi++ [121]) or in a different application domain (e.g., juvenile justice for LifeLines [87], Ph.D. education for Similan [155]). Even though results on perceptual performance and usability can to some extent be transferred to medicine, this summary focuses on studies in a medical domain with medical experts as subjects.

Controlled experiments. Solving fixed tasks in a laboratory environment allows the precise measurement of performance indicators like speed and correctness and, thus, compare different systems. Tables 5.4 and 5.5 outline nine experiments with MIVA [54, 56], IGID [28], sparklines [34], WBIVS [107], KNAVE II [91], and MTSA [101, 102].

Table 5.4. Controlled experiments comparing EHR visualization systems using medical experts as subjects.

compared systems	*task type*	*subjects*	*significant results*
MIVA (earlier, noninteractive prototype) *vs.* paper chart [56]	Explore intensive care patient status and development	16 medical residents	MIVA faster in 2 of 8 tasks MIVA more correct answers in 1 of 8 tasks
MIVA *vs.* paper chart [54]	Explore intensive care patient status and development	8 physicians, 3 nurses	No difference in correctness or speed
IGID *vs.* conventional tabular patient display [28]	Detect intensive care patient change	32 intensive care nurses	IGID detected more abnormal variables No differences in perceived workload (NASA-TLX score)
Sparklines *vs.* table [34]	Identify abnormal values and trends in lab data	12 physicians specialized in pediatric intensive care	Sparklines faster 37% unmatched interpretations across displays
WBIVS (interactive) *vs.* WBIVS (static) *vs.* tables [107]	Judge infection/rejection status of lung transplants	12 clinicians specialized in pulmonology and transplants	WBVIS (both variants) faster than table No difference in correctness
KNAVE II *vs.* ESS (electronic spreadsheet, i.e., MS Excel) *vs.* paper (MS Excel printout) [91]	Answer queries typical for oncology protocols	8 M.D./Ph.D. students, residents, and fellows	KNAVE II overall faster than printed paper; KNAVE II faster than ESS for hard and hardest queries but slower for easy queries KNAVE II overall more correct queries than paper KNAVE II most usable (SUS)
KNAVE II *vs.* ESS (electronic spreadsheet, i.e., MS Excel) [91]	Answer queries typical for oncology protocols (more complex queries)	5 physicians	KNAVE II overall faster and more correct queries KNAVE II more usable (SUS)

Table 5.5. Controlled experiments comparing EHR visualization systems using medical experts as subjects (continued).

compared systems	*task type*	*subjects*	*significant results*
MTSA (earlier prototype) *vs.* univariate time series plots with table data [101]	Predict an episode of acute hypotension in the following hour	14 internal medicine residents	No difference in correctness
MTSA *vs.* spreadsheet (MS Excel, one variable per sheet) [102]	Diagnose PDA (patent ductus arterious) in a neonate	23 pediatric residents, NICU and pediatric ICU fellows	No difference in correctness

Table 5.6. Usability tests of visualization systems using medical experts as subjects. Results generally confirm the usability of the tested systems but are mostly of a qualitative nature, which can be found in the original publications.

tested systems	*task type*	*subjects*
VisuExplore [111]	Get overview of patient development in diabetes care	9 physicians
TimeRider [122]	Explore multivariate trends in cohorts of diabetes patients	10 physicians
VIE-VISU [69]	Focus the attention toward critical situations	2 expert neonatologists
VISITORS [80]	Construct queries for patients or time intervals in oncology	5 physicians, 5 medical informaticians
VISITORS [81]	Explore database of oncology patients and use temporal abstractions	5 clinicians, 5 medical informaticians

Usability tests. Observing the interactions, insights, and difficulties of subjects, while they work with a visualization system in a laboratory environment, as well as post-test questionnaires or interviews can lead to a more general judgment of the system's usability. Table 5.6 outlines five usability tests with VisuExplore [111], TimeRider [122], VIE-VISU [69], and VISITORS [80, 81].

Case studies. These studies describe how medical experts can use the visualization system for their real work tasks. they are conducted in tight cooperation between researchers and one or a few medical experts and often run over longer periods of time, which allows the experts to become familiar with the system. Thus, case studies can give a realistic understanding of the system's strengths. Performing multiple case studies in different medical or nonmedical domains can show the generalizability [149]. Case studies were the most common approach for evaluation of EHR visualization systems in this survey. Some examples are: Lifelines2 [149], Similan [154], LifeFlow [153], EventFlow [95], OutFlow [152], VisCareTrails [88], PatternFinder [109], VISITORS [81], IPBC [45], Caregiver [36], Gravi++ [67], VisuExplore [123], Midgaard [32], or the multiple views system by Zhang et al. [159].

Deployment. An additional testimonial for the usefulness of EHR visualization is their deployment in hospitals systems or clinical research platforms—not for the sake of visualization research but for medical purposes. Four of the surveyed systems report about such deployment: ICUFiles [125] is in routine use at six intensive care units at the university hospital of Giessen, Germany. Palma et al. [104] present how the EHR system at Lucile Packard Children's Hospital at Stanford uses sparklines. Lifelines2 has been integrated in clinical research environments such as Harvard Medical School's i2b2 [97] and National Institute of Health Biomedical Translational Research Information System (BTRIS) [98] to reach thousands of researchers. The features of the VISITORS are commercially available from MediLogos [7].

Further details on the methodology and results of these studies can be found in the original papers.

5.3 Patient Data Visualization in Commercial EHR Systems

The main part of this survey was conducted based on literature search in academic publications. While this approach allows us to give an extensive view on the scientific state-of-the-art, it is not fully representative for the commercial systems actually deployed in the health care sector.

There are a number of difficulties in surveying these commercial EHR systems: Replicating the complexity of the health care sector, EHR systems need to fulfill many technical, organizational, and legal requirements. Thus, they are extensively customized for each health care institution. Often they are acquired from national or regional software providers and sometimes developed in-house by health care institutions. Thus, the number of systems that can be considered EHR systems is prohibitively high. Second, information about these systems is often not readily available to outside reviewers. Due to customizing for individual customers the details of actual deployed systems are not made public. Even if a list of capabilities is available, it rather describes the systems in general and technical terms than it elaborates on user interfaces, visualizations, or interaction. Screen shots are not always available, either. Finally, while academic publications report on the empirical evaluation of their systems, such information is usually not available from software providers. For example, it is unlikely to get access to a system's case studies or internal evaluations or obtain user population size.

Considering this, comparing systems published in scientific literature with a set of commercial systems along the same review criteria is most likely incomplete and biased. Therefore, we decided to not mix those two areas but rather include a separate section that provides a representative overview of what visualizations in commercial EHR systems are available on the market. The authors' long-lasting experience in a considerable number of research projects with collaborators from industry forms the basis of this overview. The presented systems are products we found accessible information about or one of the authors has experience with.

To avoid repetition, we structure the overview by the topics infrastructure, patient development, departmental awareness, analytics, and software libraries.

Infrastructure. EHR systems provide health care professionals not only with visualizations of patient data but also with dashboards, user interfaces for data entry, configurable alarms, and access to knowledge resources [138]. They are connected or integrated with many other systems (e.g., workflow system, laboratory, radiology, and billing) and health record exchanges. They require reliable backend databases, security, and backup.

Major vendors of such systems include Cerner [3], Epic [4], SAP [16], and Siemens [17]. Large hospital centers pioneering in electronic health record have implemented their own software and integrated their many in-hospital systems. Examples include Washington Hospital Center's Azyxxi, which later became Caradigm Amalga [2]. While most vendor systems support physicians' daily clinical tasks and administrative tasks (e.g., billing), not all support research tasks such as identifying a group of similar patients for clinical trial. The clinical tasks always take priority over the research tasks, and as such, research tasks may be supported by a completely separate system. For example, Massachusetts General Hospital's Longitudinal Medical Record (LMR) [92] is used by physicians to record all the diagnoses and treatments of patients while Research Patient Data Registry (RPDR) [105] supports the research needs of researchers.

Patient development. A frequent visualization in EHR systems is the flowsheet. This component is named after a paper-based artifact that is widely used in intensive care units. It contains key medical variables for a single patient over a period of time (e.g., 24 hours) and thus emphasizes trends and abnormal values [35, 138].

Electronic flowsheets are a combination of spreadsheets with point or line plots (e.g., Philips ICIP Critical Care [11] in Figure 5.1, Picis Critical Care Manager [13], systema mpa Fieberkurve [19] in Figure 5.2, T-Vision Fieberkurve [20] in Figure 5.3). The spreadsheet columns and the horizontal axes of plots share a common timescale. On the vertical axis, variables are displayed either as numbers in a spreadsheet row or as marks in a point or line plot. For blood pressure there is often a dedicated glyph connecting the values of systolic and diastolic blood pressure. Furthermore, there are plots for medication overview that use different icons along the horizontal time axis similar to LifeLines. Abnormal values are visually highlighted in the spreadsheets and plots.

In the flowsheet examples, the visualizations contain overlaid plots of up to five variables. These are used by conventions in some areas of medical care and often physicians expect overlaid plots in order to interpret flowsheets [94, 111]. They allow more variables to be displayed on

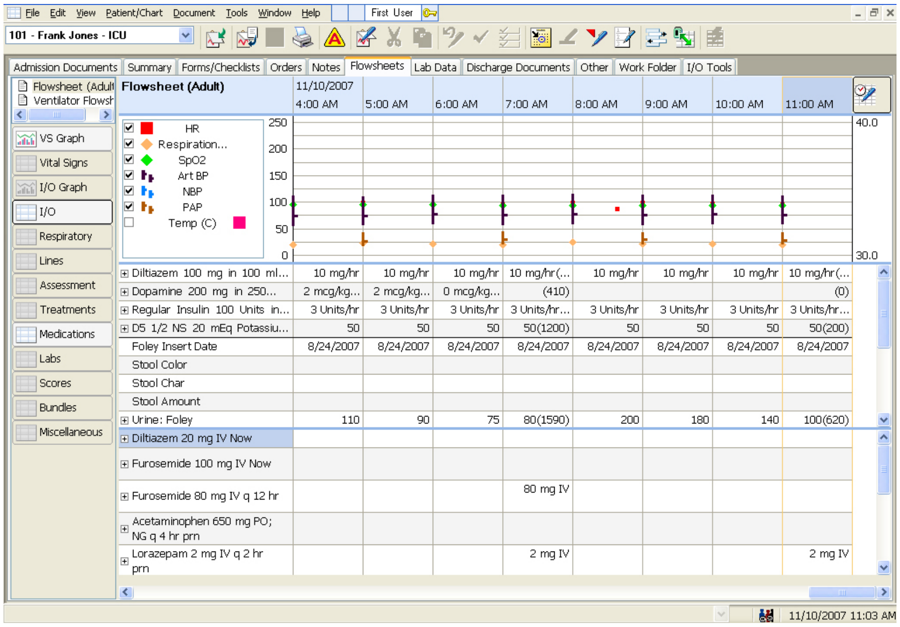

Fig. 5.1 Philips ICIP Critical Care Flowsheet with vital signs in a point plot and a spreadsheet for input/output and medication [11]. *Image courtesy of Philips Healthcare.*

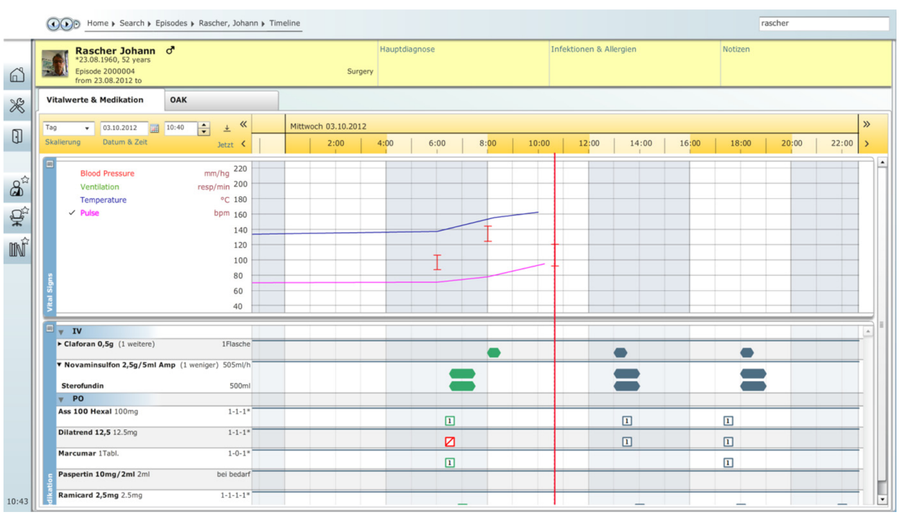

Fig. 5.2 systema mpa Fieberkurve with the "Timeline" view showing vital signs in a line plot and medication using LifeLines-based line segments and icons [19].
Image courtesy of systema Human Information Systems.

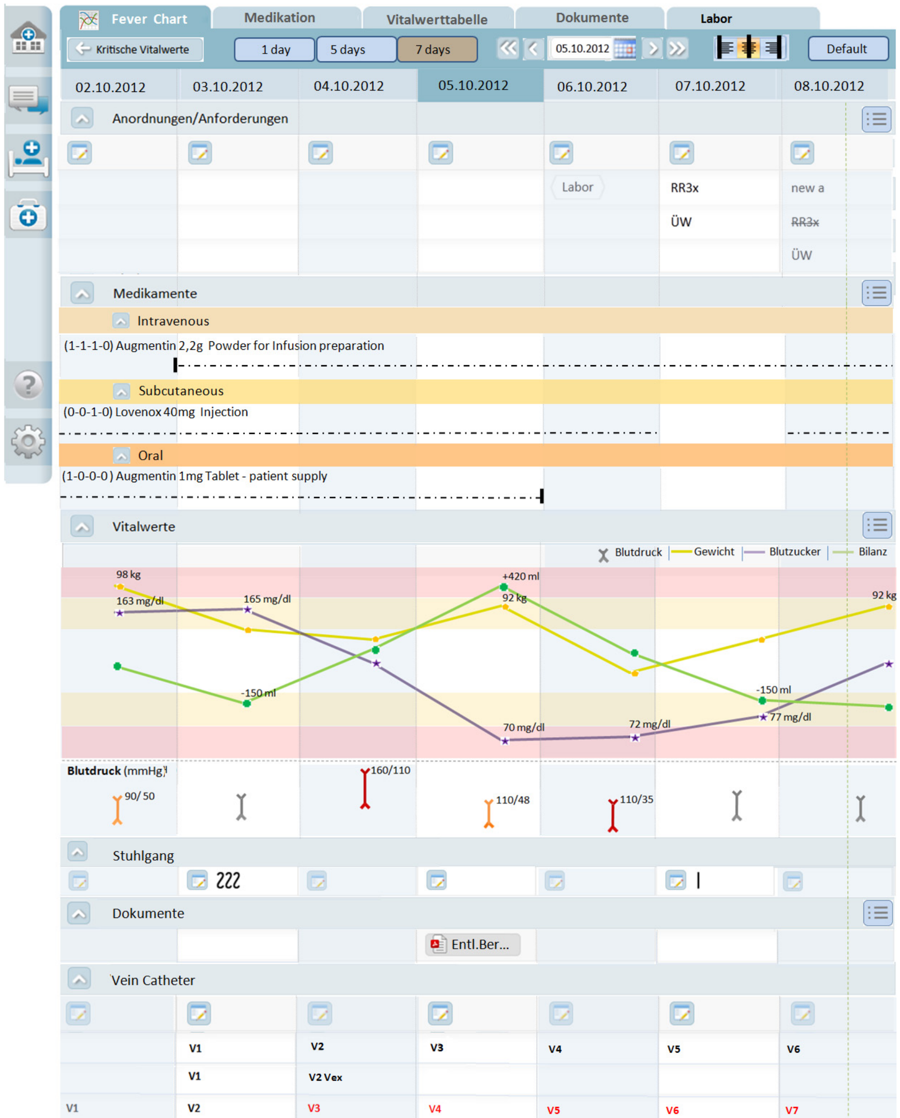

Fig. 5.3 T-Vision Fieberkurve [20] showing vital signs, medication, and other medical data over seven days. Qualitative ranges and color are used to highlight critical values. (The figure presents a user interface that was still under development at the time of writing this article.) *Image courtesy of T-Systems Austria.*

a given display area, but can have some perceptual disadvantages: first, too many marks make the plot cluttered and it becomes hard to distinct variables. Second, the variables may have diverging value ranges, so that large relative changes in a variable with a small absolute value are not visible (e.g., ventilation rate in a plot scaled for blood pressure). Using different scales for each variable makes it even harder to read the plot correctly and can imply misleading insights (cp. [94, p. 67], [146, pp. 91–94]).

Electronic flowsheets are used not only for analysis but also for data entry. They either are preconfigured or can be composed individually for clinical decision support.

Also beyond flowsheets, some interactive visualizations can be found in EHR systems. For example, Allscripts Wand Timeline (Figure 5.4) presents an interactive overview of patient data for mobile touch-screen devices. It features line plots for numerical variables and LifeLines-based line segments and icons for medication and encounters. The system offers details of these items in tooltips and can highlight items that are related to a diagnosis, which were selected from the list in top. The users can reorder panels, open or close them, and configure the presentation. They can interactively zoom and pan the time axis or highlight a date with a vertical line [1].

Departmental awareness. Overview of patients hospitalized at a department and awareness of changes in their health state can be another service of an EHR system.

For example, Picis eView [14] provides such an overview of the patient population at one or more departments or a selected subset of patients. For each patient the current value of seven numerical variables is displayed as a number along with a small bar chart that shows the trend using previous values. Both the number of the current value and bars of previous values are color-coded to highlight critical values. The details of the last six values of each variable are available in a tooltip. For even more details the user can open a patient summary page. Furthermore, an icon in the patient list notifies the user of unreviewed lab results.

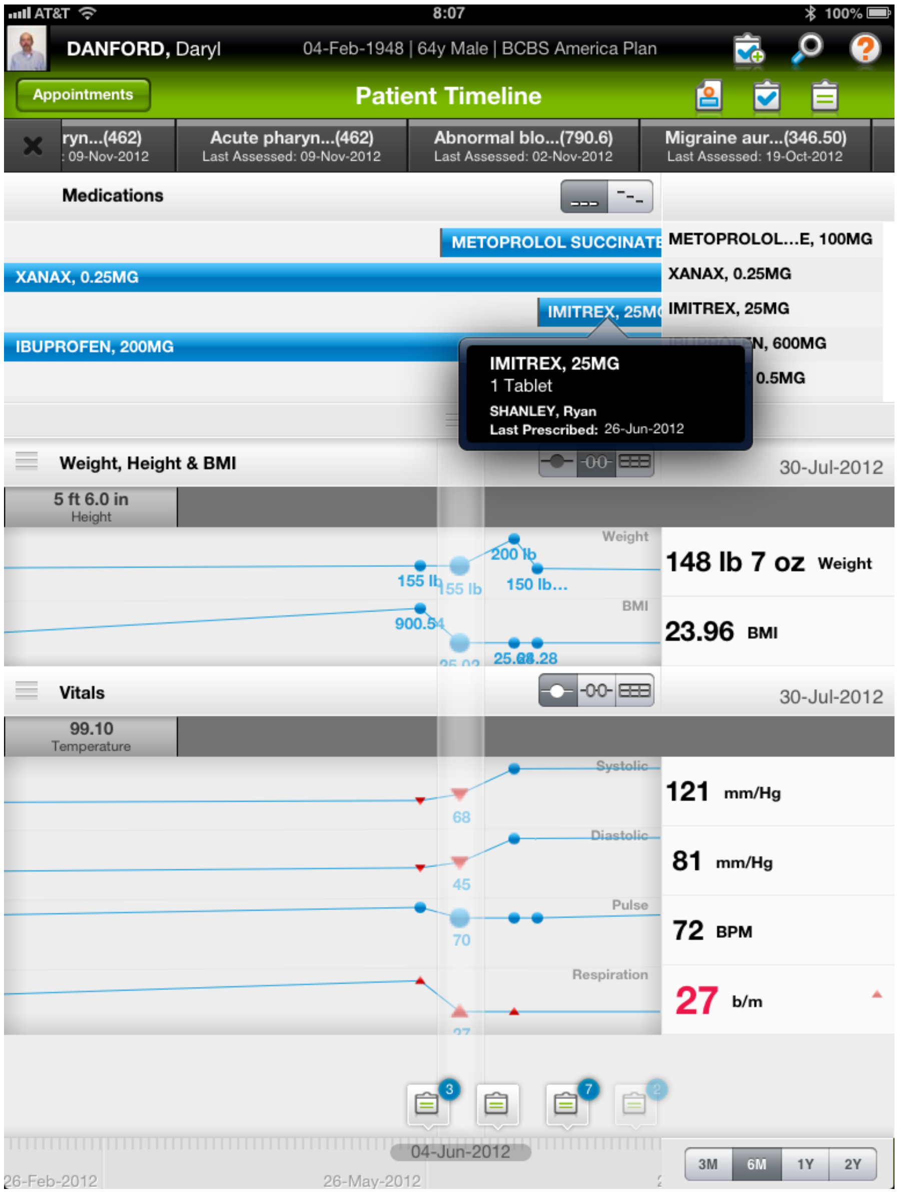

Fig. 5.4 Allscripts Wand Timeline [1] shows an overview of a patient's diagnoses, medication, vital signs, lab values, encounters along a horizontal time axis. The visualization is configurable and interactive. Here, the events of June 4, 2012 are highlighted by a vertical rule and medication details are shown in a tooltip. *Image courtesy of Allscripts.*

Analytics. EHR systems are not only used in daily clinical tasks, but also to provide feedback for quality analysis and the implementation of evidence-based care guidelines. These systems let analysts compare aggregated data, identify trends, and find outliers in a group of patients. Many systems allow multidimensional analysis of EHR data using data warehousing methods much like in business analytics systems. They often include automated analysis techniques using data mining or machine learning algorithms. IBM [5] and Oracle [9] are active in this area. However, EHR system vendors may provide this as a service, such as Caradigm [2], a collaborative effort from Microsoft and GE Healthcare.

Furthermore, the visualization methods of VISITORS are commercially available from the company MediLogos [7] and can be applied for EHR analytics.

Software libraries. Visualization techniques also exist as software widgets that can be plugged into EHR systems.

The Microsoft Health Common User Interface (CUI) [8] is an effort to standardize user interface elements in health care application across applications from different providers and, thus, improve usability and patient safety. It is comprised of extensive design guidelines and a collection of open-source software components for the .NET platform. These include a "Graphing" component for line plots and a "Timeline" component for categorical variables.

Stottler Henke DataMontage [18] is a collection of software components for interactive visualization of time-oriented data with point plots for numerical data and LifeLines visualization for nominal data. The components provide reference lines to mark clinical thresholds and have a coordinated cursor. It can be used as a Java library or an interactive software application. The software features a number of medical demos and is used to visualize EHR data in Oracle's Empirica Study On Demand product and the Pharmacovigilance Defense Application System (PVDAS) of the U.S. Army Surgeon General's office.

Summary. There is some adoption of information visualization methods in commercial systems. Apparently, the only visualizations that are in widespread use are the different implementations of flowsheets, which often use simple line plots. We did not find much information on commercial systems for visualization of multiple patients for clinical research and quality analysis but such systems are probably too specific for a specialty or user group to be advertised by major vendors.

5.4 Limitations

This survey examined in detail 14 information visualization systems that have been applied to EHRs. We also evaluated them by their support for different data types, multiple variables, multiple EHRs, and user intent. This survey was conducted based on an extensive literature search and close reading of these systems' peer-reviewed publications. When possible we contacted the authors and requested updated references and screen shots. In the previous section, we provide a perspective on visualization in commercial EHR systems. However, our process also imposed a few limitations on our survey.

In publications we have reviewed, a system's features and capabilities are often demonstrated in only one or two case studies. A more comprehensive review would include a variety of evaluation methods of each system. For example, controlled experiments could demonstrate the effectiveness of specific interaction or visualization features, while long-term case studies would demonstrate overall utility of the system. Because not every system has been evaluated in the same way and some papers do not report on the evaluation in an adequate level of detail, we have focused on the systems' features and intents instead of their overall usability or successes in real-life applications. As EHR systems mature and start incorporating visualizations we can expect more evaluation and usage reports in future surveys.

Furthermore, some of these systems are still under development, and our review relied on the descriptions found in the latest publications we could find.

5.5 Recommendations and Future Directions

Based on the EHR visualization systems presented in this survey, the obstacles we encountered in searching and comparing them, and current trends in Visualization and EHRs, the following recommendations and directions for future research can be identified.

5.5.1 Advancing Future Comparative Surveys

A general drawback of surveys is that they rely solely on the information made available by others. This survey is no exception. There is a need to promote more detailed comparison of EHR visualization systems to advance their design and to foster growth of new ideas. For the research community, the following guidelines for evaluation and result reporting would help advance this goal:

(1) Report the usage and dissemination of visualization systems in addition to their system information.
(2) Report the usage frequency of the different features of their system and their impact.
(3) Include in-depth, long-term evaluation strategies with medical professionals (in addition to controlled studies) to observe how intended end-users successfully gain insight or struggle with particular design elements [130].
(4) Apply systems to multiple case studies instead of just one or two to examine their generalizability.
(5) Give context under which intended designs succeed and contexts under which they fare less well.

One of the more difficult elements in comparing visualization systems is a lack of standard data source. The research community would greatly benefit from having a large benchmark repository of deidentified patient records with associated tasks and clinical research problems. This shared resource would encourage new research and facilitate comparison of prototypes and systems, as it happened in other communities [108], for example, the PhysioNet/Computing in Cardiology challenges in intensive care [12]. The repository can additionally be beneficial for developers of commercial EHR systems.

5.5.2 Guidelines for Information Visualization

While the systems presented here are tried and tested examples for the domain of EHRs, designers of future EHR visualization systems can also build on the body of knowledge in information visualization [40, 132, 150, 151]. It encompasses a multitude of perceptual findings. For example, Cleveland and McGill [47] empirically compared the perception of different visual mappings for numerical data and found position on a common scale to be most effective. Such findings are important for designers as they help them build interfaces that allow efficient and accurate analysis of EHR data.

Guidelines summarize this knowledge in a format that can readily be used by designers of visualization systems. The book "Show Me the Numbers" [58] is a recommendable collection of guidelines for information visualization. Specifically for health care, the Microsoft Health Common User Interface [94] contains a well-designed set of guidelines, which has largely been adopted by the British National Health Service.

5.5.3 Future Research Directions

Analyzing a single patient and the analyzing multiple patients are different goals. This difference creates a dichotomy in the visualization systems. All systems in this survey either support tasks for analyzing a single patient, or those for analyzing multiple patients. Tasks that involve both, such as comparing a single patient with multiple patients of similar history, are not extensively explored. Likewise, transitioning between a multiple patient analysis to a single patient analysis and vice versa is not widely studied.

Different tasks often require different representations of the data. For example, while using plots to show patient vital signs (for example, blood pressure) is sufficient for tracking the progress of that individual vital sign, for some higher-level analyses, the physicians only want to know whether that particular reading is abnormal. In other words, only systems that can dynamically change numerical data to and from categorical data can support tasks at different levels of abstraction. Furthermore, the increasingly large and complex data in EHRs makes it necessary to intertwine interactive visualization with automated analysis techniques as proposed by the Visual Analytics community [76]. Although VISITORS and Midgaard provide ways for creating and visualizing higher abstraction, they do not fully support richer user intents once the abstractions are created.

These first two points are important not only as standalone future directions of research, but taken together they are indicative of a larger problem. That is, visual analysis of health record is very complex. They differ in goal, in data requirement, in methodology, in scope, and in abstraction. These analyses are processes that require time, critical thinking, keen observation, and often thoughtful discussion with other knowledgeable individuals. These processes involve many tasks, of which only a few are supported in a system. There is a need to develop a well-defined process model for visual analysis of EHRs. It will allow generalization of common tasks. It will enable designers to innovate without loss of crucial functionality. More importantly, it will facilitate users to explore data thoroughly in a systematic yet flexible way [106].

Thomas and Cook [140] promote a canonical process model for analytical reasoning. Others have adapted the model for more specific domains such as social network analysis [63]. Wang et al. proposed a process model for temporal categorical data analysis [149]. Despite its narrower focus, it points out a few future opportunities this survey shares: there is a lack of support for better dissemination methods for exploratory findings, a lack of support for sharing results to promote discussion, and a lack of support for richer exploratory analysis process.

Unfortunately, building such process models is still a novel challenge. Information visualization researchers are encouraged to give input, build, and refine the evolving process model from the experience of their individual systems.

5.5.4 Alternative Target Users

The systems in this survey assume skilled professionals such as physicians, clinical researchers, or hospital quality controllers as end-users for frequent use. In this increasingly digital and self-serve world, future systems could be designed to reach patients, consumer advocates, medical policy analysts, and other intermittent users. Increasingly, patients keep their own health records and therefore have a need to review or analyze their history. In the future, a patient might present ancestral health records to physicians as part of family medical history. Web sites such as PatientsLikeMe [10] already offer such a visualization for certain chronically ill patients. Similarly, the growing interest in tracking exercise, diet, medications, and sleep patterns means that many people will need visual tools to understand daily, weekly, seasonal, and yearly patterns or changes, as is often discussed on web sites such as QuantifiedSelf.com [15].

6

Conclusions

While EHR systems make it possible for physicians or clinical researchers to retrieve important medical data on-demand and in real time, most EHR systems do not offer any advanced interaction or visualization of patient data to support clinical tasks. Containing large volumes of heterogeneous temporal data, EHRs present a variety of managerial and cognitive challenges. Real medical data are also messy, incomplete, and may contain systematic errors that impede analysis. We believe that effective information visualization techniques can both reduce the problems medical analysts face and facilitate their analysis tasks.

This work surveys state-of-the-art information visualization systems found in the scientific literature. These systems have all been applied to medical domains, but offer different visualization and interaction techniques for corresponding analytical tasks. Designers of future user interfaces for EHR systems will find this survey useful to familiarize themselves with the features and virtues of existing work. Thus, they become even more capable to tackle the challenges ahead.

Acknowledgments

We thank Danube University Krems, Vienna University of Technology, systema Human Information Systems, NÖ Landeskliniken-Holding, Landesklinikum Krems, NÖGUS, Washington Hospital Center, and Oracle Health Sciences for their support. The work presented here was partially supported by the National Institute of Health (grant RC1CA147489-02: Interactive Exploration of Temporal Patterns in Electronic Health Records), Grant No. 10510592 for Patient-Centered Cognitive Support under the Strategic Health IT Advanced Research Projects Program (SHARP) from the Office of the National Coordinator for Health Information Technology, the FWF Austrian Science Fund (project number: P22883), and the Laura Bassi Centre of Expertise CVAST (project number: 822746).

Furthermore, we thank Dominique Brodbeck, Carlo Combi, Anthony Faiola, Theresia Gschwandtner, Werner Horn, Denis Klimov, Patricia Ordóñez Rozo, David Pieczkiewicz, and Art Swanson for providing us with feedback and screenshots of their systems, as well as Todd Johnson, Paolo Buono, James Cimino, Paolo Federico, Seth Powsner, and the anonymous reviewers for their feedback on the manuscript. We appreciate the guidance and support from editor-in-chief Ben Bederson.

References

[1] "Allscripts Wand," http://allscripts.com/en/solutions/ambulatory-add-ons/wand.html (cited 19 Nov 2012).

[2] "Caradigm Amalga," http://www.caradigm.com/pages/en-us/products/amalga.aspx (cited 20 Nov 2012).

[3] "Cerner Acute Care Electronic Medical Record," http://www.cerner.com/solutions/Hospitals_and_Health_Systems/Acute_Care_EMR/ (cited 7 Nov 2012).

[4] "EpicCare Inpatient Clinical System," http://www.epic.com/software-inpatient.php (cited 20 Nov 2012).

[5] "IBM Healthcare industry solutions," http://www.ibm.com/industries/healthcare/ (cited 20 Nov 2012).

[6] "Interactive Visual Exploration of Electronic Health Records," HCIL 2008 Symposium Workshop, University of Maryland, http://www.cs.umd.edu/hcil/ehrviz-workshop/ (cited 24 Jul 2012), June 2008.

[7] "MediLogos Data Visualization products," https://sites.google.com/a/medilogos.com/medilogos/products/data-visualization (cited 7 Nov 2012).

[8] "Microsoft Health Common User Interface," http://mscui.codeplex.com/ (cited 7 Nov 2012).

[9] "Oracle Health Sciences," http://www.oracle.com/us/industries/health-sciences/overview/ (cited 20 Nov 2012).

[10] "PatientsLikeMe," http://www.patientslikeme.com/ (cited 8 Aug 2012).

[11] "Philips IntelliVue Clinical Information Portfolio (ICIP) Critical Care," http://www.healthcare.philips.com/pwc_hc/main/shared/Assets/Documents/patient_monitoring/icip/ICIP_Concept_WP_10.pdf (cited 7 Nov 2012).

[12] "PhysioNet/Computing in Cardiology Challenges," http://physionet.org/challenge/ (cited 3 Aug 2012).

[13] "Picis Critical Care Manager," http://www.picis.com/solutions/critical-care/critical-care-manager.aspx (cited 7 Nov 2012).

[14] "Picis eView," http://www.picis.com/solutions/critical-care/eview-for-critical-care-manager.aspx (cited 7 Nov 2012).

[15] "QuantifiedSelf portal," http://www.quantifiedself.com/ (cited 3 Aug 2012).

[16] "SAP for Healthcare," http://www.sap.com/industries/healthcare/index.epx (cited 20 Nov 2012).

[17] "Siemens Healthcare Hospital IT," http://healthcare.siemens.com/hospital-it (cited 20 Nov 2012).

[18] "Stottler Henke DataMontage," http://www.stottlerhenke.com/datamontage/ (cited 7 Nov 2012).

[19] "systema mpa Fieberkurve," http://www.systema.info/uploads/media/mpa_CGM_G3_SYSTEMA_Fieberkurve.pdf (cited 7 Nov 2012).

[20] "T-Systems Österreich Lösungen Health," http://www.t-systems.at/loesungen/healthcare/481368 (cited 20 Nov 2012).

[21] "Visual Analytics in Healthcare," Workshop in conjunction with VisWeek 2011, Oct 23-28, Providence, RI, http://www.visualanalyticshealthcare.org/VAHC_2011/ (cited 24 Jul 2012).

[22] C. Ahlberg, C. Williamson, and B. Shneiderman, "Dynamic queries for information exploration: An implementation and evaluation," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*, pp. 619–626, 1992.

[23] W. Aigner, K. Kaiser, and S. Miksch, "Visualization techniques to support authoring, execution, and maintenance of clinical guidelines," in *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*, (A. ten Teije, P. Lucas, and S. Miksch, eds.), pp. 140–159, Amsterdam: IOS, 2008.

[24] W. Aigner and S. Miksch, "CareVis: Integrated visualization of computerized protocols and temporal patient data," *Artificial Intelligence in Medicine*, vol. 37, no. 3, pp. 203–218, July 2006.

[25] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski, "Visual methods for analyzing time-oriented data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 47–60, 2008.

[26] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of Time-Oriented Data*. London: Springer, 2011.

[27] W. Aigner, A. Rind, and S. Hoffmann, "Comparative evaluation of an interactive time-series visualization that combines quantitative data with qualitative abstractions," *Computer Graphics Forum*, vol. 31, no. 3, pp. 995–1004, 2012.

[28] S. Anders, R. Albert, A. Miller, M. B. Weinger, A. K. Doig, M. Behrens, and J. Agutter, "Evaluation of an integrated graphical display to promote acute change detection in ICU patients," *International Journal of Medical Informatics*, vol. 81, no. 12, pp. 842–851, December 2012.

[29] P. André, M. L. Wilson, A. Russell, D. A. Smith, A. Owens, and m.c. schraefel, "Continuum: Designing timelines for hierarchies, relationships and scale," in *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST'07)*, pp. 101–110, October 2007.

[30] C. Ardito, P. Buono, M. F. Costabile, and R. Lanzilotti, "Two different interfaces to visualize patient histories on a PDA," in *Proceedings of the Conference on Human-Computer Interaction with Mobile Devices and Services*, pp. 37–40, Helsinki, Finland: ACM, 2006.

[31] J. S. Ash, M. Berg, and E. Coiera, "Some unintended consequences of information technology in health care: The nature of patient care information system-related errors," *Journal of the American Medical Informatics Association*, vol. 11, no. 2, pp. 104–112, 2004.

[32] R. Bade, 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*, (E. Dykstra-Erickson and M. Tscheligi, eds.), pp. 105–112, Vienna, Austria: ACM, 2004.

[33] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis, "Algorithms for drawing graphs: An annotated bibliography," *Computational Geometry*, vol. 4, no. 5, pp. 235–282, October 1994.

[34] D. T. Bauer, S. Guerlain, and P. J. Brown, "The design and evaluation of a graphical display for laboratory data," *Journal of the American Medical Informatics Association*, vol. 17, no. 4, pp. 416–424, July 2010.

[35] D. T. Bauer, S. A. Guerlain, and P. J. Brown, "Evaluating the use of flowsheets in pediatric intensive care to inform design," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, no. 10, pp. 1054–1058, October 2006.

[36] D. Brodbeck, R. Gasser, M. Degen, S. Reichlin, and J. Luthiger, "Enabling large-scale telemedical disease management through interactive visualization," in *Connecting Medical Informatics and Bio-Informatics, Proceedings of the International Congress Eur. Fed. Med. Inform. (MIE)*, pp. 1172–1177, 2005.

[37] A. A. T. Bui, D. R. Aberle, and H. Kangarloo, "TimeLine: Visualizing integrated patient records," *IEEE Transactions on Information Technology in Biomedicine*, vol. 11, no. 4, pp. 462–473, July 2007.

[38] P. Buono, A. Aris, C. Plaisant, A. Khella, and B. Shneiderman, "Interactive pattern search in time series," in *Proceedings of the Conference Visualization and Data Analysis (VDA 2005)*, pp. 175–186, SPIE, March 2005.

[39] P. Buono, C. Plaisant, A. Simeone, A. Aris, B. Shneiderman, G. Shmueli, and W. Jank, "Similarity-based forecasting with simultaneous previews: A river plot interface for time series forecasting," in *Proceedings of the International Conference Information Visualization (IV '07)*, pp. 191–196, July 2007.

[40] S. K. Card, J. D. Mackinlay, and B. Shneiderman, eds., *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann, 1999.

[41] S. Carpendale, "Evaluating information visualizations," in *Information Visualization: Human-Centered Issues and Perspectives*, (A. Kerren, J. T. Stasko, J. Fekete, and C. North, eds.), pp. 19–45, Berlin: Springer, 2008.

[42] L. Chittaro, "Information visualization and its application to medicine," *Artificial Intelligence in Medicine*, vol. 22, no. 2, pp. 81–88, 2001.

[43] L. Chittaro, "Visualization of patient data at different temporal granularities on mobile devices," in *Proceedings of the Working Conference Advanced Visual Interfaces (AVI'06)*, pp. 484–487, Venezia, Italy: ACM, May 2006.

[44] L. Chittaro and C. Combi, "Visualizing queries on databases of temporal histories: New metaphors and their evaluation," *Data & Knowledge Engineering*, vol. 44, no. 2, pp. 239–264, February 2003.

[45] L. Chittaro, C. Combi, and G. Trapasso, "Data mining on temporal data: A visual approach and its clinical application to hemodialysis," *Journal of Visual Languages and Computing*, vol. 14, no. 6, pp. 591–620, December 2003.

[46] T. Christensen and A. Grimsmo, "Instant availability of patient records, but diminished availability of patient information: A multi-method study of GP's use of electronic patient records," *BMC Medical Informatics and Decision Making*, vol. 8, no. 1, p. 12, March 2008.

[47] W. S. Cleveland and R. McGill, "Graphical perception: theory, experimentation, and application to the development of graphical methods," *Journal American Statistical Association*, vol. 79, no. 387, pp. 531–554, September 1984.

[48] A. Cockburn, A. Karlson, and B. B. Bederson, "A review of Overview+Detail, zooming, and Focus+Context interfaces," *ACM Computing Surveys*, vol. 41, no. 1, pp. 2:1–2:31, December 2008.

[49] C. Combi, E. Keravnou-Papailiou, and Y. Shahar, *Temporal Information Systems in Medicine*. New York: Springer, 2010.

[50] C. Combi and B. Oliboni, "Visually defining and querying consistent multigranular clinical temporal abstractions," *Artificial Intelligence in Medicine*, vol. 54, no. 2, pp. 75–101, February 2012.

[51] C. Combi, L. Portoni, and F. Pinciroli, "Visualizing temporal clinical data on the WWW," in *Proceedings of the Artificial Intelligence in Medicine, AIMDM'99*, (W. Horn, Y. Shahar, G. Lindberg, S. Andreassen, and J. Wyatt, eds.), pp. 301–311, Springer, January 1999.

[52] S. B. Cousins and M. G. Kahn, "The visual display of temporal information," *Artificial Intelligence in Medicine*, vol. 3, no. 6, pp. 341–357, December 1991.

[53] J. A. Fails, A. Karlson, L. Shahamat, and B. Shneiderman, "A visual interface for multivariate temporal data: Finding patterns of events across multiple histories," in *Proceedings of the IEEE Symposium on Visual Analytics Science And Technology (VAST)*, pp. 167–174, Baltimore, MD, 2006.

[54] A. Faiola, C. Boston-Clay, J. Jones, M. Downey, and C. M. Newlon, "Managing patient health across diverse spaces: Using activity theory to model pervasive decision support," in *Proceedings of the CHI2012 Workshop "Bridging Clinical and Non-Clinical Health Practice: Opportunities and Challenges"*, 2012. http://www.ics.uci.edu/~yunanc/chi2012health/Downloads/Faiola.pdf (cited 15 Aug 2012).

[55] A. Faiola and S. Hillier, "Multivariate relational visualization of complex clinical datasets in a critical care setting: A data visualization interactive prototype," in *Proceedings of the International Conference Information Visualization (IV 2006)*, pp. 460–468, July 2006.

[56] A. Faiola and C. Newlon, "Advancing critical care in the ICU: A Human-Centered biomedical data visualization systems," in *Ergonomics and Health Aspects, Proceedings of the International Conference on Human-Computer Interaction (HCII 2011)*, (M. M. Robertson, ed.), pp. 119–128, Springer, 2011.

[57] G. Falkman, "Information visualisation in clinical odontology: Multidimensional analysis and interactive data exploration," *Artificial Intelligence in Medicine*, vol. 22, no. 2, pp. 133–158, May 2001.

[58] S. Few, *Show Me the Numbers: Designing Tables and Graphs to Enlighten*. Oakland, CA: Analytics Press, Second ed., 2012.

[59] T. Fonseca, C. Ribeiro, and C. Granja, "Vital signs in intensive care: Automatic acquisition and consolidation into electronic patient records," *Journal of Medical Systems*, vol. 33, no. 1, pp. 47–57, May 2008.

[60] G. W. Furnas, "Generalized fisheye views," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '86)*, pp. 16–23, 1986.

[61] D. Gotz, J. Sun, N. Cao, and S. Ebadollahi, "Visual cluster analysis in support of clinical decision intelligence," in *Proceedings of the AMIA 2011 Symposium*, pp. 481–490, 2011.

[62] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang, "CareCruiser: Exploring and visualizing plans, events, and effects interactively," in *Proceedings of the IEEE Pacific Visualization Symposium (PacificVis)*, pp. 43–50, 2011.

[63] D. L. Hansen, D. Rotman, E. Bonsignore, N. Milic-Frayling, E. M. Rodrigues, M. Smith, and B. Shneiderman, "Do you know the way to SNA?: A process model for analyzing and visualizing social media data," Technical Report HCIL-2009-17, University of Maryland, http://hcil2.cs.umd.edu/trs/2009-17/2009-17.pdf (cited 21 Aug 2012), 2009.

[64] R. L. Harris, *Information Graphics: A Comprehensive Illustrated Reference*. New York: Oxford University Press, 1999.

[65] C. F. Herot, "Spatial management of data," *ACM Transactions on Database Systems*, vol. 5, no. 4, pp. 493–513, December 1980.

[66] D. U. Himmelstein, A. Wright, and S. Woolhandler, "Hospital computing and the costs and quality of care: A national study," *American Journal of Medicine*, vol. 123, no. 1, pp. 40–46, January 2010.

[67] K. Hinum, S. Miksch, W. Aigner, S. Ohmann, C. Popow, M. Pohl, and M. Rester, "Gravi++: Interactive information visualization to explore highly structured temporal data," *Journal of Universal Computer Science*, vol. 11, no. 11, pp. 1792–1805, November 2005.

[68] H. Hochheiser and B. Shneiderman, "Dynamic query tools for time series data sets: Timebox widgets for interactive exploration," *Information Visualization*, vol. 3, no. 1, pp. 1–18, 2004.

[69] W. Horn, C. Popow, and L. Unterasinger, "Support for fast comprehension of ICU data: Visualization using metaphor graphics," *Methods of Information in Medicine*, vol. 40, no. 5, pp. 421–424, 2001.

[70] Human-Computer Interaction Lab, University of Maryland, "EventFlow: Exploring point and interval event temporal patterns," http://www.cs.umd.edu/hcil/eventflow (cited 29 Dec 2012).

[71] A. Inselberg and B. Dimsdale, "Parallel coordinates: A tool for visualizing multi-dimensional geometry," in *Proceedings of the Conference on Visualization (Vis'90)*, pp. 361–378, 1990.

[72] Institute of Medicine, Committee on Patient Safety and Health Information Technology, *Health IT and Patient Safety: Building Safer Systems for Better Care*. Washington, D.C.: National Academies, 2012. http://www.nap.edu/catalog.php?record_id=13269 (cited 1 Aug 2012).

[73] D. F. Jerding and J. T. Stasko, "The information mural: A technique for displaying and navigating large information spaces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 4, no. 3, pp. 257–271, 1998.

[74] J. Jin and P. Szekely, "Interactive querying of temporal data using a comic strip metaphor," in *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (VAST 2010)*, pp. 163–170, October 2010.

[75] H. C. Karadimas, C. Chailloleau, F. Hemery, J. Simonnet, and E. Lepage, "Arden/J: An architecture for MLM execution on the java platform," *Journal of the American Medical Informatics Association*, vol. 9, pp. 359–368, 2002.

[76] D. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann, eds., *Mastering The Information Age — Solving Problems with Visual Analytics*. Goslar, Germany: Eurographics, 2010. http://www.vismaster.eu/book/ (cited 9 Aug 2012).

[77] D. A. Keim, H. Kriegel, and T. Seidl, "Visual feedback in querying large databases," in *Proceedings of the IEEE Conf. Visualization*, pp. 158–165, 1993.

[78] P. R. Keller and M. M. Keller, *Visual Cues: Practical Data Visualization*. Los Alamitos: IEEE Computer Society, 1993.

[79] R. Kincaid and H. Lam, "Line graph explorer: Scalable display of line graphs using focus+context," in *Proceedings of the Working Conference on Advanced Visual Interfaces*, pp. 404–411, 2006.

[80] D. Klimov, Y. Shahar, and M. Taieb-Maimon, "Intelligent selection and retrieval of multiple time-oriented records," *Journal of Intelligent Information Systems*, vol. 35, pp. 261–300, 2010.

[81] D. Klimov, Y. Shahar, and M. Taieb-Maimon, "Intelligent visualization and exploration of time-oriented data of multiple patients," *Artificial Intelligence in Medicine*, vol. 49, no. 1, pp. 11–31, 2010.

[82] R. Koppel, J. P. Metlay, A. Cohen, B. Abaluck, A. R. Localio, S. E. Kimmel, and B. L. Strom, "Role of computerized physician order entry systems in facilitating medication errors," *JAMA*, vol. 293, no. 10, pp. 1197–1203, March 2005.

[83] R. Kosara and S. Miksch, "Visualization methods for data analysis and planning in medical applications," *International Journal of Medical Informatics*, vol. 68, no. 1–3, pp. 141–153, December 2002.

[84] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale, "Empirical Studies in Information Visualization: Seven Scenarios," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 9, pp. 1520–1536, September 2012.

[85] T. Lammarsch, W. Aigner, A. Bertone, J. Gärtner, E. Mayr, S. Miksch, and M. Smuc, "Hierarchical temporal patterns and interactive aggregated views for pixel-based visualizations," in *Proceedings of the International Conference Information Visualisation (IV 2009)*, pp. 44–50, 2009.

[86] B. J. Lesselroth and D. S. Pieczkiewicz, "Data visualization strategies for the electronic health record," in *Advances in Medicine and Biology,* vol. 16, (L. V. Berhardt, ed.), pp. 107–140, NOVA, 2011.

[87] D. Lindwarm Alonso, A. Rose, C. Plaisant, and K. L. Norman, "Viewing personal history records: A comparison of tabular format and graphical presentation using LifeLines," *Behaviour & Information Technology*, vol. 17, no. 5, pp. 249–262, September 1998.

[88] L. Lins, M. Heilbrun, J. Freire, and C. Silva, "VisCareTrails: Visualizing trails in the electronic health record with timed word trees, a pancreas cancer use case," in *Proceedings of the IEEE VisWeek Workshop on Visual Analytics in Health Care*, (J. J. Caban and D. Gotz, eds.), pp. 13–16, 2011.

[89] M. Lungu and K. Xu, "Biomedical information visualization," in *Human-Centered Visualization Environments*, (A. Kerren, A. Ebert, and J. Meyer, eds.), pp. 311–342, Berlin: Springer, 2007.

[90] T. Mabotuwana, J. Warren, and J. Kennelly, "A computational framework to identify patients with poor adherence to blood pressure lowering medication," *International Journal of Medical Informatics*, vol. 78, no. 11, pp. 745–756, November 2009.

[91] S. B. Martins, Y. Shahar, D. Goren-Bar, M. Galperin, H. Kaizer, L. V. Basso, D. McNaughton, and M. K. Goldstein, "Evaluation of an architecture for intelligent query and exploration of time-oriented clinical data," *Artificial Intelligence in Medicine*, vol. 43, no. 1, pp. 17–34, May 2008.

[92] Massachusetts General Hospital, "Longitudinal Medical Record (LMR) tutorial," http://www.massgeneral.org/stoecklecenter/programs/patient_exper/lmr.aspx (cited 29 Nov 2012).

[93] P. McLachlan, T. Munzner, E. Koutsofios, and S. North, "LiveRAC: Interactive visual exploration of system management time-series data," in *Proceedings of the Annual SIGCHI Conference on Human Factors in Computing Systems*, pp. 1483–1492, 2008.

[94] Microsoft, "Design guidance — Displaying graphs and tables," Microsoft Health Common User Interface, http://www.mscui.com/DesignGuide/Pdfs/Design%20Guidance%20--%20Displaying%20Graphs%20and%20Tables.pdf (cited 21 Aug 2012), 2008.

[95] M. Monroe, R. Lan, J. M. del Olmo, B. Shneiderman, C. Plaisant, and J. Millstein, "The intervals and absences of temporal query," in *Proceedings of the 2013 Annual Conference Human Factors in Computing Systems (CHI '13)*, 2013. forthcoming.

[96] T. Munzner, "Visualization principles," Keynote at Workshop on Visualizing Biological Data (VIZBI 2011), http://www.cs.ubc.ca/~tmm/talks/vizbi11/vizbi11.pdf (cited 24 Jul 2012), March 2011.

[97] S. N. Murphy, M. Mendis, K. Hackett, R. Kuttan, W. Pan, L. C. Phillips, V. Gainer, D. Berkowicz, J. P. Glaser, I. Kohane, and H. C. Chueh, "Architecture of the Open-source Clinical Research Chart from Informatics for Integrating Biology and the Bedside," in *Proceedings of the AMIA 2007 Symposium*, pp. 548–552, 2007.

[98] National Institutes of Health, "Biomedical translational research information system," http://btris.nih.gov/ (cited 9 Aug 2012).

[99] M. J. O'Connor, S. W. Tu, and M. A. Musen, "The Chronus II temporal database mediator," in *Proceedings of the AMIA 2002 Annual Symposium*, pp. 567–571, 2002.

[100] P. Ordóñez, M. desJardins, C. Feltes, C. U. Lehmann, and J. Fackler, "Visualizing multivariate time series data to detect specific medical conditions," in *Proceedings of the AMIA 2008 Symposium*, pp. 530–534, 2008.

[101] P. Ordóñez, M. desJardins, M. Lombardi, C. U. Lehmann, and J. Fackler, "An Animated Multivariate Visualization for Physiological and Clinical Data in the ICU," in *Proceedings of the 1st ACM International Health Informatics Symposium (IHI 2010)*, pp. 771–779, 2010.

[102] P. Ordonez, T. Oates, M. E. Lombardi, G. Hernandez, K. W. Holmes, J. Fackler, and C. U. Lehmann, "Visualization of multivariate time-series data in a neonatal ICU," *IBM Journal of Research and Development*, vol. 56, no. 5, pp. 7:1–7:12, October 2012.

[103] J. M. Overhage, B. Mamlin, J. Warvel, J. Warvel, W. Tierney, and C. J. McDonald, "A tool for provider interaction during patient care: G-CARE.," in *Proceedings of the Annual Symposium on Computer Application in Medical Care*, pp. 178–182, 1995.

[104] J. P. Palma, P. J. Brown, C. U. Lehmann, and C. A. Longhurst, "Neonatal informatics: Optimizing clinical data entry and display," *NeoReviews*, vol. 13, no. 2, pp. 81–85, February 2012.

[105] Partners Research Computing, "Research Patient Data Registry (RPDR)," http://rc.partners.org/rpdr (cited 29 Nov 2012).

[106] A. Perer and B. Shneiderman, "Systematic yet flexible discovery: guiding domain experts through exploratory data analysis," in *Proceedings of the International Conference on Intelligent User Interfaces (IUI '08)*, pp. 109–118, ACM, 2008.

[107] D. 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 *Proceedings of the AMIA 2007 Symposium*, pp. 598–602, 2007.

[108] C. Plaisant, J. Fekete, and G. Grinstein, "Promoting Insight-Based evaluation of visualizations: From contest to benchmark repository," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 120–134, 2008.

[109] C. Plaisant, S. Lam, B. Shneiderman, M. S. Smith, D. Roseman, G. Marchand, M. Gillam, C. Feied, J. Handler, and H. Rappaport, "Searching electronic health records for temporal patterns in patient histories: A case study with Microsoft Amalga," in *Proceedings of the AMIA 2008 Symposium*, pp. 601–605, AMIA, November 2008.

[110] 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*, pp. 76–80, 1998.

[111] M. Pohl, S. Wiltner, A. Rind, W. Aigner, S. Miksch, T. Turic, and F. Drexler, "Patient development at a glance: An evaluation of a medical data visualization," in *Proceedings of the IFIP Human-Computer Interaction — INTERACT 2011, Part IV*, (P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler, eds.), pp. 292–299, Springer, 2011.

[112] A. R. Post and J. H. Harrison Jr., "PROTEMPA: A method for specifying and identifying temporal sequences in retrospective data for patient selection," *Journal of the American Medical Informatics Association*, vol. 14, no. 5, pp. 674–683, September 2007.

[113] G. Potamias, "State of the art on systems for data analysis, information retrieval and decision support," INFOBIOMED project, Deliverable D13, http://www.infobiomed.org/paginas_en/INFOBIOMED_D13_final.pdf (cited 1 Aug 2012), 2006.

[114] S. M. Powsner and E. R. Tufte, "Graphical summary of patient status," *Lancet*, vol. 344, no. 8919, pp. 386–389, August 1994.

[115] S. M. Powsner and E. R. Tufte, "Summarizing clinical psychiatric data," *Psychiatric Services*, vol. 48, pp. 1458–1461, November 1997.

[116] B. Preim and D. Bartz, *Visualization in Medicine: Theory, Algorithms, and Applications.* Amsterdam: Morgan Kaufmann, 2007.

[117] B. Purin, M. Rigo, S. Riccadonna, and S. Forti, "An interactive real-time visualization environment for patients with heart failure," in *Proceedings of the AMIA 2005 Symposium*, p. 1088, 2005.

[118] R. Rao and S. K. Card, "The table lens: Merging graphical and symbolic representations in an interactive focus + context visualization for tabular information," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating interdependence*, pp. 318–322, ACM, 1994.

[119] L. Rasmussen and J. Starren, "Augmented interactive starfield displays for health management," in *Proceedings of the AMIA 2007 Symposium*, p. 1087, 2007.

[120] H. Reijner, "The development of the horizon graph," in *Proceedings of the Vis08 Workshop "From Theory to Practice: Design, Vision and Visualization"*, (L. Bartram, M. Stone, and D. Gromala, eds.), 2008.

[121] M. Rester, M. Pohl, S. Wiltner, K. Hinum, S. Miksch, C. Popow, and S. Ohmann, "Mixing evaluation methods for assessing the utility of an interactive InfoVis technique," in *Proceedings of the International Conference Human-Computer Interaction (HCII 2007), Part I*, (J. Jacko, ed.), pp. 604–613, Springer, July 2007.

[122] A. Rind, W. Aigner, S. Miksch, S. Wiltner, M. Pohl, F. Drexler, B. Neubauer, and N. Suchy, "Visually exploring multivariate trends in patient cohorts using animated scatter plots," in *Ergonomics and Health Aspects of Work with Computers, Proceedings of the International Conference held as part of HCI International 2011*, (M. M. Robertson, ed.), pp. 139–148, Springer, 2011.

[123] A. Rind, W. Aigner, S. Miksch, S. Wiltner, M. Pohl, T. Turic, and F. Drexler, "Visual exploration of time-oriented patient data for chronic diseases: Design study and evaluation," in *Information Quality in e-Health, Proceedings of the USAB 2011*, (A. Holzinger and K. Simonic, eds.), pp. 301–320, 2011.

[124] G. Robertson, R. Fernandez, D. Fisher, B. Lee, and J. Stasko, "Effectiveness of animation in trend visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1325–1332, 2008.

[125] R. Röhrig, M. Benson, A. Michel-Backofen, D. Brammen, F. Brenck, K. Marquardt, and G. Hempelmann, "Designing a graphical user interface for an intensive-care unit patient data chart," in *Proceedings of the International Congress of the European Federal Medical Informatics (MIE)*, pp. 1055–1060, 2005.

[126] T. Saito, H. N. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, and T. Kaseda, "Two-tone Pseudo coloring: Compact visualization for one-dimensional data," in *Proceedings of the 2005 IEEE Symposium Information Visualization (InfoVis 2005)*, pp. 173–180, 2005.

[127] Y. Shahar and C. Cheng, "Intelligent visualization and exploration of time-oriented clinical data," in *Proceedings of the Annual Hawaii International Conference on System Sciences (HICSS-32)*, 1999. doi:10.1109/HICSS.1999.773019.

[128] Y. Shahar, D. Goren-Bar, D. Boaz, and G. Tahan, "Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions," *Artificial Intelligence in Medicine*, vol. 38, no. 2, pp. 115–135, October 2006.

[129] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," in *Proceedings of th IEEE Symposium Visual Languages*, pp. 336–343, 1996.

[130] B. Shneiderman and C. Plaisant, "Strategies for evaluating information visualization tools: Multi-dimensional in-depth long-term case studies," in *Proceedings of the 2006 AVI Workshop BEyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV '06)*, pp. 38–43, ACM, 2006.

[131] R. Smith, "What clinical information do doctors need?," *BMJ*, vol. 313, no. 7064, pp. 1062–1068, October 1996.

[132] R. Spence, *Information Visualization: Design for Interaction*. Upper Saddle River, NJ: Prentice Hall, Second ed., 2007.

[133] M. Spenke, "Visualization and interactive analysis of blood parameters with InfoZoom," *Artificial Intelligence in Medicine*, vol. 22, no. 2, pp. 159–172, May 2001.

[134] W. W. Stead and H. S. Lin, eds., *Computational Technology for Effective Health Care: Immediate Steps and Strategic Directions*. Washington, D.C.: National Academies Press, 2009.

[135] M. Steenwijk, J. Milles, M. A. van Buchem, J. H. Reiber, and C. P. Botha, "Integrated visual analysis for heterogeneous datasets in cohort studies," in *IEEE VisWeek Workshop on Visual Analytics in Health Care*, (D. Gotz and J. Sun, eds.), 2010. http://research.ihost.com/vahc2010/pdfs/steenwijk.pdf (cited 2 Aug 2012).

[136] S. S. Stevens, "On the theory of scales of measurement," *Science*, vol. 103, no. 2684, pp. 677–680, 1946.

[137] E. Sundvall, M. Nystrom, M. Forss, R. Chen, H. Petersson, and H. Ahlfeldt, "Graphical overview and navigation of electronic health records in a prototyping environment using google earth and openEHR archetypes," in *Proceedings of the 12th World Congress of the Health (Medical) Informatics (Medinfo 2007)*, (K. A. Kuhn, J. R. Warren, and T. Leong, eds.), pp. 1043–1047, Amsterdam, 2007.

[138] P. C. Tang and C. J. McDonald, "Electronic health record systems," in *Biomedical Informatics*, (E. H. Shortliffe and J. J. Cimino, eds.), pp. 447–475, New York: Springer, 2006.

[139] C. Tao, K. Wongsuphasawat, K. Clark, C. Plaisant, B. Shneiderman, and C. G. Chute, "Towards event sequence representation, reasoning and visualization for EHR data," in *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium (IHI '12)*, pp. 801–806, 2012.

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

[141] M. Tory and T. Möller, "Rethinking visualization: A high-level taxonomy," in *Proceedings of the IEEE Symp. Information Visualization (INFOVIS 2004)*, pp. 151–158, 2004.

[142] E. R. Tufte, *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.

[143] E. R. Tufte, *Beautiful Evidence*. Cheshire, CT: Graphics Press, 2006.

[144] J. W. Tukey, *Exploratory Data Analysis*. Reading, MA: Addison-Wesley, 1977.

[145] L. Tweedie, B. Spence, D. Williams, and R. Bhogal, "The attribute explorer," in *Conference Companion on Human Factors in Computing Systems (CHI '94)*, pp. 435–436, ACM, 1994.

[146] H. Wainer, *Visual Revelations: Graphical Tales of Fate and Deception from Napoleon Bonaparte to Ross Perot*. New York, NY: Copernicus, 1997.

[147] 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 Annual SIGCHI Conference on Human Factors in Computing Systems*, pp. 457–466, ACM, 2008.

[148] T. D. Wang, C. Plaisant, B. Shneiderman, N. Spring, D. Roseman, G. Marchand, V. Mukherjee, and M. Smith, "Temporal summaries: Supporting temporal categorical searching, aggregation and comparison," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 6, pp. 1049–1056, 2009.

[149] T. D. Wang, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman, "Extracting insights from electronic health records: Case studies, a visual analytics process model, and design recommendations," *Journal of Medical Systems*, vol. 35, no. 5, pp. 1135–1152, May 2011.

[150] M. Ward, G. G. Grinstein, and D. Keim, *Interactive Data Visualization: Foundations, Techniques, and Applications*. A K Peters, May 2010.

[151] C. Ware, *Information Visualization: Perception for Design*. San Francisco, CA: Morgan Kaufmann, Second ed., 2004.

[152] K. Wongsuphasawat and D. H. Gotz, "Outflow: Visualizing patient flow by symptoms and outcome," in *Proceedings of the IEEE VisWeek Workshop on Visual Analytics in Health Care*, pp. 25–28, 2011.

[153] 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 2011 Annuals Conference on Human Factors in Computing Systems (CHI '11)*, pp. 1747–1756, 2011.

[154] K. Wongsuphasawat, C. Plaisant, M. Taieb-Maimon, and B. Shneiderman, "Querying event sequences by exact match or similarity search: Design and empirical evaluation," *Interacting with Computers*, vol. 24, no. 2, pp. 55–68, March 2012.

[155] K. Wongsuphasawat and B. Shneiderman, "Finding comparable temporal categorical records: A similarity measure with an interactive visualization," in *Proceedings of the IEEE Symposium Visual Analytics Science and Technology 2009 (VAST 2009)*, pp. 27–34, IEEE, 2009.

[156] M. Workman, M. F. Lesser, and J. Kim, "An exploratory study of cognitive load in diagnosing patient conditions," *International Journal for Quality in Health Care*, vol. 19, no. 3, pp. 127–133, June 2007.

[157] J. S. Yi, Y. ah Kang, J. T. Stasko, and J. A. Jacko, "Toward a deeper understanding of the role of interaction in information visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1224–1231, 2007.

[158] Z. Zhang, F. Ahmed, A. Mittal, I. V. Ramakrishnan, R. Zhao, A. Viccellio, and K. Mueller, "AnamneVis: A framework for the visualization of patient history and medical diagnostics chains," in *Proceedings of the IEEE VisWeek Workshop on Visual Analytics in Health Care*, (J. J. Caban and D. Gotz, eds.), pp. 17–20, 2011.

[159] Z. Zhang, A. Mittal, S. Garg, A. E. Dimitriyadi, I. V. Ramakrishnan, R. Zhao, A. Viccellio, and K. Mueller, "A visual analytics framework for emergency room clinical encounters," in *IEEE Workshop on Visual Analytics in Health Care*, 2010. http://research.ihost.com/vahc2010/pdfs/zhang.pdf (cited 2 Aug 2012).