AnamneVis: A Framework for the Visualization of Patient History and Medical Diagnostics Chains

Zhiyuan Zhang[1], Faisal Ahmed[1], Arunesh Mittal[1], IV Ramakrishnan[1], Rong Zhao[1], Asa Viccellio[2], and Klaus Mueller[1]

[1]Computer Science Department and Center for Wireless and Information Technology (CEWIT)

[2]Department of Emergency Medicine

Stony Brook University

ABSTRACT

The medical history or *anamnesis* of a patient is the factual information obtained by a physician for the medical diagnostics of a patient. This information includes current symptoms, history of present illness, previous treatments, available data, current medications, past history, family history, and others. Based on this information the physician follows through a medical diagnostics chain that includes requests for further data, diagnosis, treatment, follow-up, and eventually a report of treatment outcome. Patients often have rather complex medical histories, and visualization and visual analytics can offer large benefits for the navigation and reasoning with this information. Here we present *AnamneVis*, a system where the patient is represented as a radial sunburst visualization that captures all health conditions of the past and present to serve as a quick overview to the interrogating physician. The patient's body is represented as a stylized body map that can be zoomed into for further anatomical detail. On the other hand, the reasoning chain is represented as a multi-stage flow chart, composed of date, symptom, data, diagnosis, treatment, and outcome.

KEYWORDS: health care, medical record presentation, EHR, EMR

1 INTRODUCTION

The electronic health record (EHR) digitally stores patient health information generated by one or more clinical encounters in any care delivery setting. This information includes patient demographics, problems, symptoms, diagnoses, progress notes, treatments, medication, vital signs, past medical history, immunizations, laboratory data, radiology reports, and many others. However, the acceptance of the EHR in clinical practice lags far behind its expectation and potential. Related information and overviews are typically difficult to obtain, severely impeding a physician's diagnostic reasoning. The inefficient, fragmented display of patient information is a likely cause. In this paper we offer a first step to overcome these deficiencies by comprehensibly organizing the patient medical history, also known as *anamnesis*. We employ the concept of *Five W's* (*who*, *when*, *what*, *where*, *wh*y, and also *how*) of journalistic reporting to structure the medical information domain and provide a suitable visual mapping for each for visual information display.

The Five W's are the elements of information needed to get a full story. They are encountered in many playing fields: by a journalist uncovering a political scandal, a police detective investigating a crime, a customer service representative trying to resolve a complaint, and a market analyst planning an effective marketing campaign. The order in which the information is gathered or interrogated can vary case by case – crucial is only that all five W's are ultimately addressed.

When it comes to applying the Five W's to visualization design, we can break it down into two steps: (1) identify all Five W components and their relations, and (2) map these to suitable visual information encodings and interactions.

We propose to use the Five W's in our health care informatics application as a means to establish a comprehensive multi-faceted assessment of the patient and his (her) history for intuitive information retrieval. The goal is information organization and integration along these various aspects. Overview and detail-on-demand requires hierarchies, and effective information organization requires robust encoding by ways of well-established criteria – we use standard codes commonly used for billing in hospitals which enables us to easily build our system on top of an existing health care information system. These codes are ICD, CPT, and NDC. ICD is the code used to describe the condition or disease being treated, also known as the diagnosis. CPT is the code used to describe medical services and procedures performed by doctors for a particular diagnosis. NDC is the code used for administered drugs. ICD is widely accessible (developed by the World Health Organization), CPT is proprietary and only available to healthcare providers, and NDC is also publicly available. Further goals, often expressed by our collaborating emergency physician – who is also a co-author of this article – are ease of information access and flexibility in displayed aggregated information and data. To enable this functionality, our system is fully interactive and the displays are fully linked and coordinated.

2 RELATED WORK

A number of approaches for the visualization of medical patient records have been proposed, and new systems are likely to emerge as the Electronic Health Record (EHR) is adopted widely. A frequent paradigm is to organize the patient records along the time axis. Prominent efforts in that direction are LifeLines [7] and LifeLines2 [11] in which health records are distinguished by their inherent aspects, such as problems, symptoms, tests/results, diagnosis, treatments and medications, etc. and color is used to indicate severity or type A level of detail mechanism allows one to zoom into patient records. A number of other works, such as [6], have also embraced this type of patient data visualization. Particularly interesting in this context is the work of Aigner et al. [2] who have made use of illustrative abstractions to gradually transition between broad qualitative overviews of temporal data (for example, blood pressure) to detailed, quantitative time signals. These techniques are part of the Midgaard system [3] which also provides a visualization scheme in which acquired patient data are mapped to a template of a human body (although little further detail on how this scheme is used in practice is available). The system described in [8] gathers close-ups of acquired radiological data around a volume-rendered full body. In fact, many modern EHR systems now support time-line views and are also beginning to support body-centric data layouts. Another frequently used paradigm is that of flow-charts, as used in clinical
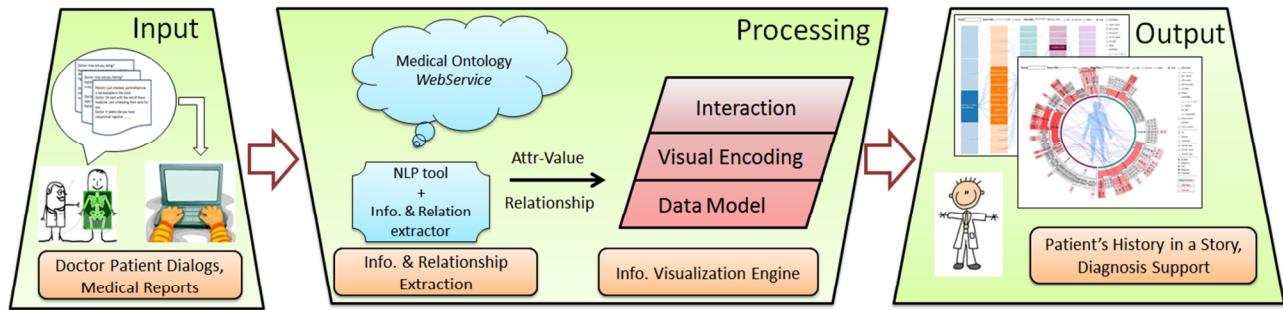
Figure 1: System Pipeline

algorithm maps [5] and others [4][10], where patient records are visualized as a logical execution sequence of plans. These methods typically operate without temporal alignments. Finally, works also exist that combine these two paradigms into coordinated views [1].

While our recent work [12] also embraced the Five W's scheme, its main focus was a visual interface that a doctor might use to log and review evidence gathered (and actions required) during a patient-doctor dialog (called *encounter*). This system combined the temporal functionality of [7][11] with the body-centric data arrangement of [3][8] and supported analytical reasoning with these information items via a force-directed graph (called the *diagnosis sandbox*).

When cast into the Five W's we find that most existing systems support the *when*, *what*, *why*, and *where* aspects quite well, although few support all of these. Functionality for coordinated views linking specialized visualization for these aspects is less supported. Apart from this, a further main difference to existing systems is our representation of the *who*. While most systems reduce it to simple personal data, such as name, age, gender, smoker, and the like, we see it as an opportunity to represent all medical information ever recorded about a patient – a true reflection of the person (in terms of medical history at least). All is captured within a modern information visualization framework and linked with the other coordinated displays for the other 4 W aspects.

3   IDENTIFYING THE FIVE W'S: INFORMATION EXTRACTION

The information flow of our system is summarized in Fig. 1. The input to our system are patient records and medical reports, doctor-patient dialogs and other input, results from triage, and data acquired from the patient, such as radiological images, lab analyses, and the like. At the processing stage an NLP (Natural Language Processing) engine cooperates with an online medical ontology server to extract structured information and relationships from this incoming information and data stream. It then formats the extraction results into the Five W model and passes it on to the visualization engine. The visualization engine has all procedures and data models to encode the Five W information facets into corresponding visuals and interaction procedures. The output of this process is then presented in the visual interface that is subject of this paper. In the box labelled 'Output' we show two displays: (i) a hierarchical radial ring display (foreground window) that visualizes the patient history in the context of a centered body map and (ii) a sequential (causal) display (background window) that visualizes the diagnostic reasoning chain. Before we describe this interface in detail, we first discuss the conceptual information organization of our system, in terms of the structuring Five W's.

3.1   The *Who* and *What*

The *who* and *what* information helps doctors to quickly assess the history and status of the patient. It describes the patient in terms of:

♦ Symptoms and Diagnosis: this includes the patient's symptoms, injuries, and any diagnosed diseases. All of this information can be encoded using the ICD code standard.

♦ Procedures: these include patient tests and examinations, treatments administered, and drugs prescribed. This type of information can be encoded using the CPT code or the ICD-procedure code standard, and the NCD code standard.

♦ Data: these include test and examination results, review of systems, vital signs, and social and family history. The codes for these are part of the procedure code and yields information on what the patient already has.

Our system encodes this information in two ways: in a hierarchical radial ring display and in a sequential (causal) display.

3.2   The *Where*

The *where* information refers to the location of the *who* and *what* information within the confines of the patient's body. While not all information can be localized that way, for the information that can be localized, we encode it in a body outline map surrounded by the ring display. Items on the ring display are pointing to the appropriate locations on the body. The Google Body Browser [13] could then be indexed by a subset of the *what* and so give the doctors a good start for further exploration and also offer explanations to the patient.

3.3   The *When*, *Why*, and *How*

The *when*, *why*, and *how* show a case under (doctor) collaborative diagnosis/treatment, or an entire life span. It demonstrates for each node *what, when, why*, and *how* that node appears. Various multi-resolution and selection techniques are available to make the visualization scalable. It supports two types of displays: a sequential display and a hierarchical radial display.

The sequential display stresses causal relationships and encourages causal reasoning done by the doctor. It also aims to model the typical medical workflow: (1) observe symptoms and possibly browse history data, (2) prescribe and evaluate tests results, (3) form hypotheses and possibly acquire more data, (4) cast diagnoses and (5) prescribe treatments. These steps may all be executed within one patient visit or they may prolong over some period of time, but the overall workflow is always engaged. The 5th step may include a referral to another doctor, which then starts another workflow (back-linking to the previous).

The hierarchical radial display aims to provide an overview of the entire history of the patient, offering detail on demand.

4 ENCODING THE FIVE W'S: INFORMATION VISUALIZATION

We have two types of cooperating displays:

♦ A hierarchical radial (patient overview) display with an integrated body outline primarily for the *who* and *where*.

♦ A sequential (diagnostic reasoning) display primarily for the *when*, *why*, and *how*.

The *what* is part of both displays (in form of the various nodes) and is context-sensitive. The two interfaces are linked, such that operations on either view will be reflected in the other. Thus, one can quickly switch between the sequential (and possibly evolving) diagnostic reason flow and the radial patient overview display.

4.1 Hierarchical Radial Display – The 'Who' Display

The hierarchical radial display is used primarily to show the *who* and *where* information of the patient. Based on the category information discussed in Section 3.1, the *who* includes three radial displays, one for symptoms and diagnoses, one for procedures and treatments, and one for data. These three displays are interlinked to allow doctors to obtain a full picture of the patient as well as assess existing relationships.

4.1.1 Data Model

We use a tree data structure to store the code hierarchy information. For each symptom or diagnosis the patient has, we find the node *n* in the tree with the corresponding ICD9 (soon ICD10) code, and insert the new item as a child for node *n*. For example, if the patient has *bacterial meningitis* whose ICD9 code is 320, we first build an incident node *m* for this diagnosis to store its information (severity, result, etc.). In the tree, we find the node *n* with code 320, which is [320 *bacterial meningitis*]. Then we insert *m* as a child of *n*. After this, we update all ancestors of *n* with the new inserted incident node's information, such as number of incidents that fall into this category, severity, and so on. By doing this for all symptoms, diagnosis, and procedures, the tree will always be current and contain the patient's entire history.

4.1.2 Visual Design

We use the *sunburst* visualization paradigm [9] (see Figure 2) to visually convey the tree structure. A sunburst is a polar-coordinates hierarchical space-filling diagram. Nodes in the sunburst layout are drawn as solid areas (either wedges or bars), and their placement relative to adjacent nodes reveals the

Figure 1: Node design. Color encodes severity. The main node layer tells us that the patient has a relative severe disease in the nervous and sense organs. The children layers provide more detail with regards to what the diseases are.

relationships in the hierarchy. Because the nodes are space-filling the angle for each node can be used to encode additional information, such as number of incidents in our case.

Each node has a wedged shape in the sunburst tree. We further decompose the node into three layers to encode more information, as is shown in Fig. 1. These layers are:

♦ *Layer 1:* the main node layer which is used to display information about the node, such as codes, name, etc.

♦ *Layer 2:* encodes the next lower level in the hierarchy. It is meant to give users a quick overview on the sub-diseases without showing their real nodes. We provide this layer to make the hierarchy display scalable.

Each radial display is either hierarchy-centric or patient-centric. In the hierarchy-centric display (Fig. 2a), each node in the sunburst tree is sized by how many sub-categories it has. It focuses more on the hierarchy information represented in the medical codes and serves as an illustration of the complexity of a sub-system and its composition. In the patient-centric display on the other hand, more radial space is dedicated for diagnoses/procedures the patient had activities in. For categories that the patient does not have any activities in, the node will be collapsed to save space for others (see Fig. 2b).

There are three levels of code hierarchies in the sunburst radial display. The first level corresponds to the highest code hierarchy level. The second level shows more detailed categories. The third level contains the incident nodes, which are the medical items (symptoms/procedures/diagnosis) that the patient has activities in. Three default level filters are provided to help users quickly explore these three levels. Also users can expand and collapse the nodes interactively by their expertise.

The root of the tree is displayed in the center of the sunburst. However given the sole application context – the patient – we chose to replace the standard root node by a body outline. This enables us to intuitively fuse the *who* with the *where* display. If an incident (medical record) has corresponding location information,
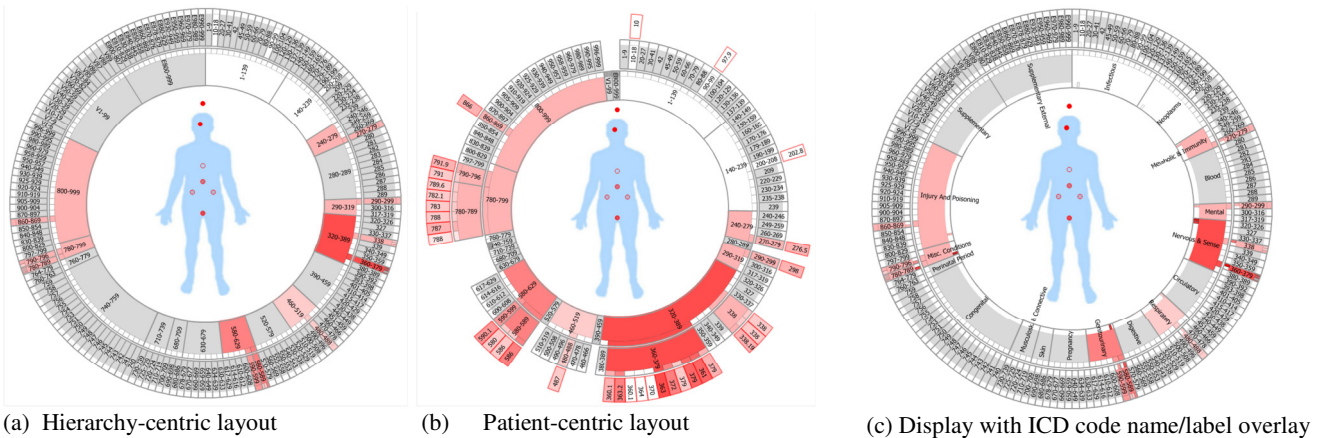

(a) Hierarchy-centric layout

(b) Patient-centric layout

(c) Display with ICD code name/label overlay

Figure 2: Sunburst display for symptoms and diagnoses

3

a red dot is displayed in the body outline. The intensity is used to encode the severity. Thus by looking at the body outline, doctors can quickly learn which parts of the patient's body have (or had) diseases and also judge their severity by the color intensity. Hovering on the red dots will popup more details about the injured part, such as name, severity, and how many incidents are related. Clicking the red dot will highlight the corresponding diseases in the sunburst tree.

Finally, users unfamiliar with the ICD9 coding system have the option to display the ICD node names and labels as an overlay (see Fig. 2c). Since the nodes are always in the same relative positions in the sunburst display, experienced users may soon acquire a mental map of the system and only require the overlay in non-routine situations.

4.2 Sequential (Causal) Display

The sequential display (Fig. 3) is used mainly to demonstrate the *what*, *when*, *why* and *how* information, which embodies the medical diagnostic flow. The medical records are organized by an underlying graph data structure. Each node corresponds to one incident (medical primitive), which could be a doctor visit, symptom, test/data, diagnosis or treatment. Edges represent relationships.

4.2.1 Visual Design

A node is displayed as one elongated box because it better utilizes the rectangular screen, better fits the text, and has better scalability compared to a circular shape. All of our medical collaborators agreed on this. If two nodes are related with one another, an edge is drawn to link them together. Edge bundling is used to reduce cluttering. Usually the diagnostic workflow is: Patient visits doctor → patient complains about symptoms → doctor orders tests for patient → doctor renders a diagnosis → treatments are given → outcome is observed. Thus, the sequential display can show these reasoning chains very well. In some cases the current doctor refers the patient to see another specialist (which is the treatment in this case), or current symptoms are caused by previous described drugs (which can be a form of diagnosis). In situations back edges appear. Back edges are shown in different color (red) to make them easy to see. Back edges may be due to treatments causing new symptoms, or they may be treatments constituting doctor referrals. Fig. 4 has no back edges.

5 CONCLUSIONS

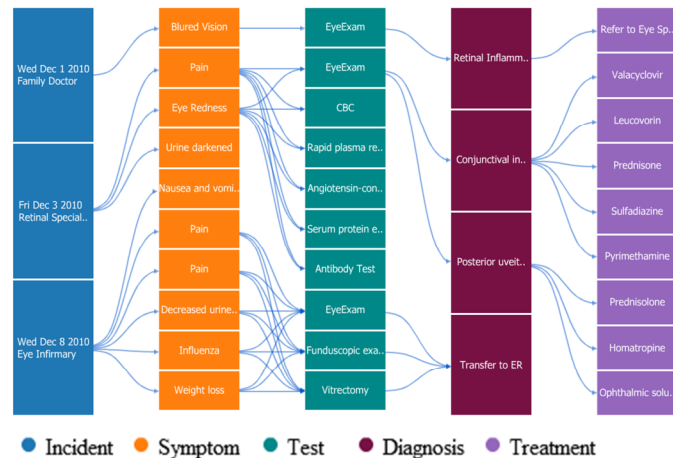We have presented an application of the Five W's scheme of information gathering and reporting, with a special application to health care informatics. While our more informal user studies are highly encouraging and promising, we would like to conduct more formal studies next. We have collaborations with six leading physicians at our university, ranging from radiology, neurosurgery, family medicine, and emergency medicine. Using well-defined tasks, we would now like to test our system with these individuals and also with the good-sized population of medical students they educate. Finally, a second possible user group for our framework are medical coding personnel who work in the hospital billing office to translate medical records to ICD9 code. Our sunburst display has good potential for them to better recognize relationships in medical services and so perform more accurate billing statements. We are currently pursuing efforts on this level as well.

REFERENCES

[1] W. Aigner, S. Miksch, "Supporting Protocol-Based Care in Medicine via Multiple Coordinated Views," *Proc. Coordinated and Multiple Views in Exploratory Visualization*, pp. 118-129, 2004.

[2] W. Aigner, S. Miksch, W. Müller, H. Schumann, C. Tominski, "Visual Methods for Analyzing Time-Oriented Data," *IEEE Trans. on Visualization and Computer Graphics,* 14(1):47-60, 2008.

[3] R. Bade, S. Schlechtweg, S. Miksch, "Connecting Time-oriented Data and Information to a Coherent Interactive Visualization," *Proc. Human Factors in Computing Systems (CHI),* pp. 105-112, 2004.

[4] J. Fox and R. Thomson., "Decision Support and Disease Management: A Logic Engineering Approach," *IEEE Transactions on Information Technology in Biomedicine*, 2(4):217–228, 1998.

[5] D. Hadorn, "Use of Algorithms in Clinical Practice Guideline Development: Methodology Perspectives," *AHCPR Pub.*, 0009(95):93–104, Jan. 1995.

[6] R. Kosara, S. Miksch, "Visualization Techniques for Time-Oriented, Skeletal Plans in Medical Therapy Planning," *Proc. Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making (AIMDM),* pp. 291-300, 1999.

[7] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, B. Shneiderman, "Lifelines: Using visualization to enhance navigation and analysis of patient records," *Proc. AMIA Annual Symposium*, pp. 76–80, 1998.

[8] T. Ropinski, I. Viola, M. Bierman, H. Hauser, K. Hinrichs, "Multimodal Visualization with Interactive Closeups," EGUK Theory and Practice of Computer Graphics (TPCG), 2009.

[9] J. Stasko and E. Zhang. Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. *IEEE Symposium on Information Visualization, 2000. InfoVis 2000*, pages 57-65, 2000.

[10] S. Quaglini, M. Stefanelli, G. Lanzola, V. Caporusso, S. Panzarasa, "Flexible guideline-based patient careflow systems," *Artificial Intelligence in Medicine*, 22(1):65–80, 2001.

[11] T. Wang, C. Plaisant, A. Quinn, R. Stanchak, B. Shneiderman, and S. Murphy, "Aligning temporal data by sentinel events: Discovering patterns in electronic health records," *Proc. ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 457-466, 2008.

[12] Z. Zhang, A. Mittal, S. Garg, A.Dimitriyadi, IV Ramakrishnan, R. Zhao, A. Viccellio, K. Mueller, "A Visual Analytics Framework for Emergency Room Clinical Encounters," *IEEE VisWeek Workshop on Visual Analytics in Health Care*, 2010.

[13] Google Body Browser: http://bodybrowser.googlelabs.com

Figure 4. Sequential display for diagnostic chain.