A Visual Analytics Framework for Emergency Room Clinical Encounters

Zhiyuan Zhang[1], Arunesh Mittal[1], Supriya Garg[1], Alexander E. Dimitriyadi[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

High costs, lack of speed, non-intuitive interfaces, and inefficient, fragmented display of patient information have hindered the adoption of the Electronic Health Record (EHR). Critical factors inhibiting adoption of the EMR include the time spent by the health care providers (HCP) in accessing and also documenting patient information during clinical encounters. We describe an emerging visual analytics system dedicated to clinical encounters in emergency room scenarios. It unifies all EMR information fragments into a single interactive visual framework, controlled by voice and touch, in which physicians can conduct diagnostic reasoning tasks in a direct data and information centric manner. We illustrate our system by ways of a typical clinical scenario and point out directions for future research and development.

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

1 INTRODUCTION

The medical record is an ancient concept, dating back to the days of Hippocrates, the father of Western medicine (around 400 BC). He prescribed the medical record as an entity that (1) should accurately reflect the course of disease and (2) should indicate the probable cause of disease [1]. These two fundamental goals still very much stand, but are now enhanced by the additional functionalities enabled by the electronic processing, storage, and transmission capabilities of the *electronic* medical record (EMR) or more appropriately named electronic health record (EHR). The EHR digitally stores patient health information generated by one or more clinical encounters in any care delivery setting, such as patient demographics, problems, symptoms, diagnoses, progress notes, treatments, medications, vital signs, past medical history, immunizations, laboratory data, radiology reports, etc. Thus, referring back at the two rules set forth by Hippocrates, the new, electronic form of the medical record now enables better record keeping and mobility (advancing rule #1) and also more refined and advanced reasoning with regards to the possible cause of the disease (advancing rule #2). Important is also that the EHR promises to promote evidence-based medicine as an acceleration of the growth in the knowledge of modern medicine. While these are arguably good features, they do come at a considerable price, which also deflates claims that the EHR saves cost. With EHRs the clinician spends much time with mundane documentation tasks entering data into myriads of electronic forms – time he/she then lacks for the patient care itself. Furthermore, current systems for analytical exploration of clinical data and information are still far from mature, leading to further waste of expensive clinician time, and also reducing the benefit for the patient.

As a goal, the ideal EHR system should be intuitive and quick, integrating and displaying comprehensive information in an easily digestible and meaningful format, support and enhance clinical decision-making, and allow for simple, rapid input and assimilation of data from a variety of sources, including input from the health care provider (HCP). The ability to directly and immediately act on information then minimizes the need for multitasking, and allows the HCP to focus on issues of higher complexity. As such, the ideal EMR should have a self-explanatory, uncluttered user interface which presents information relevant to clinical decision making, and thus require minimal training. Furthermore, it should fully integrate actions with their documentation. Thus, as a bottom-line, it should save the HCP time and help guide care.

In this paper, we report on our progress towards these goals, describing a clinician-focused system which combines state of the art information visualization technologies with voice and multi-touch interaction capabilities as well as data analytics. The structure of our paper is as follows. Section 2 provides a brief overview on related work in this area of research. Section 3 presents our overall vision. Section 4 describes our system in detail using several case studies to illustrate the system features. Finally, Section 5 ends our paper with conclusions and pointers to future work.

2 RELATED WORK

A number of approaches for the visualization of medical patient records have been proposed. A frequent paradigm is to organize the patient records along the time axis. Very active efforts in that direction are LifeLines [11] and now LifeLines2 [14] in which 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 then allows one to zoom into patient records. A number of other works, such as [10], 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 [12] gathers close-ups of acquired radiological data around a volume-rendered full body. In fact, many modern EMR 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 algorithm maps [8] and others [6][13], 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].

Our proposed framework uses some elements of these existing techniques, but adds new functionalities to them, and in addition, also introduces a number of metaphors, methods, and information organization principles that have, to the best of our knowledge, not been described so far.

3 INFORMATION MANAGEMENT IN THE EMERGENCY ROOM

Emergency room (ER) physicians must deal with multifarious pieces of information in a fast-paced clinical environment. In this section, we first give some insight into the complexity of the information management task an ER physician typically faces in daily clinical routine, and then present the vision, written from the perspective of an ER physician (Dr. Asa Viccellio, one of the co-authors of this paper) that has been serving as a guideline in the ongoing development of our system.

3.1 The complexity of information in the ER

The health care provider must contribute to the EMR by documenting various tasks outlined below; perform the crucial task of integrating disparate pieces of information; verify the accuracy of each; find associations and contradictions; note changes in the patient's status and react appropriately; and keep track of and document key elements and actions. These tasks must be simultaneously accomplished for multiple patients at once. For instance, the following would represent a typical information task on a patient with fever in the ER:

1. Obtain a history of the present illness
2. Review and add to past history, medications, family history, social history, and review of systems, keeping mentally "on the table" those elements possibly relevant to the new presentation
3. Perform a physical examination and keeping "on the table" key findings
4. Order laboratory data; know at a given moment in time what is ready and what is not; note and track new abnormalities, compare current results with prior results, and note differences; note and "shelve" those things which are normal
5. Obtain radiographic studies and EKG, with processes similar to 4
6. Request and obtain consultation from appropriate services as needed, tracking events from initiation of request to completion of consultation
7. Respond to new information against the background of what is (or isn't) known about the patient, and take appropriate action within an appropriate time course
8. Reevaluate and readjust therapy and diagnosis as circumstances warrant
9. Keep track of elements 1-8 while caring for multiple other patients

3.2 Interacting with information in the ER: a vision

With most current systems, the "table" is the physician's mental space. The physician must maintain an internal mental array, from multiple sources, of relevant information on multiple patients simultaneously, to remember to note and act on findings that arise during the encounter. Some actions will be immediate, and some may be delayed well beyond the current visit. In typical current EMRs, the action typically is unlinked to an actionable finding in the sense that the finding does not directly trigger an action (although this begun to change). The physician conceives of the action at the bedside, but in general must cease current activity, log into a system, call up the patient's record, and activate the ordering module to implement the action. In the midst of this, the physician might be diverted to another task and fail to accomplish the currently conceived task.

To illustrate how these interactions should be done more ideally let us assume the following fictitious ER clinical encounter. John Smith is a cancer patient who has come to the emergency department. He has been triaged and placed in a room. The physician enters the room, and the screen by the patient is activated. As he does so, all available patient information important to the physician appears on a large touch screen adjacent to the patient's bedside, with the ability to drill down to specific detail, call up an x-ray or EKG, and compare prior results. Having this information immediately available will assist the physician in understanding the patient's background in relationship to the current visit, minimize duplication of work. The patient can help the clinician verity the accuracy of the information displayed (but the screen could always be inactivated if the patient requests it, or as the clinician determines necessary). The interview begins. If the physician wishes to document at the bedside during the clinical encounter, the physician taps on the documentation icon. The typical sections of an interview are History, Past History, Family History, Social History, Review of Systems, and Physical Exam. If a nurse or resident has already entered information in these sections, the physician can review and verify this information with the patient, and modify/add to any part of the record as needed. With the selection of any of these sections (by voice or touch), voice dictation with voice recognition is initiated. Voice recognition appears on the screen, if desired, and the physician can approve it, ignore it for the time being, or send it to the ER clerk/ transcriptionist for correction. If new information is obtained out of sequence, tapping on that section and initiating voice dictation will append the dictation to prior dictation for that section. At any time, the history can be reactivated and new information added. When reviewing the history, any part of the documentation can be selected by dragging a finger across the words; dragging this selection to the appropriate spot on the screen will add this phrase to the active problem list, to minimize the chance of forgetting a key diagnosis or management issue. Verification of transcription can occur in real time, or be completed after revision by transcription personnel. New information derived from text analysis will appear on the touch screen for confirmation by the clinician. Although all of this can occur away from the bedside, the least time is spent by real time bedside documentation.

The exam may proceed as follows. "Let me examine you," the physician states. The physician exam menu appears on the screen, with a body diagram. Either by voice or tapping screen, the physician can activate either general voice dictation for the examination, or body specific areas, to enhance voice recognition. The physician can tap on the head of the body diagram, and say "tenderness and swelling"; if needed, the physician could also photograph this area, and the photo will be uploaded to the EMR and linked to the head exam. When the exam is complete, a body graphic is displayed. The system then extracts data from the examination and generates a visual representation on screen: all normal areas examined and shown in light green, abnormal areas in dark red and unexamined areas in white.

The "orders" icon is tapped or, alternatively, the physician states "we're going to see why you have a fever", and a standard set of labs, x-rays, EKG, and antipyretics are presented. Also, tapping the temperature icon will pop up the same information. All of these specific orders, as well as other optional orders, appear on the screen. By tapping, voice, or gesture, any specific order can be toggled on or off, or modified. If the set is acceptable as is, a simple gesture, or “OK”, completes the orders. Because of the swollen head, the physician also states "Let’s get a head CT to see what is going on."

This order is entered into the system, and is received by the radiology department, along with the reason for the test. Throughout this process, there has been no logging on and off; no typing; no need to leave the room to initiate a care plan for this patient. The physician can now leave and proceed directly to the next patient in need of care. At any subsequent time, when the physician enters the room, new results and pending results appear on the screen. The physician will be alerted to any serious abnormalities already via a page alert.

Similarly, when the nurse enters, patient information and pending orders appear on the screen. An IV must be started and bloods must be drawn. Labels for the blood specimens have been printed out in the room. While assembling materials, the nurse taps the "VS" icon, and vital signs are automatically obtained and downloaded into the EMR.

After placing the IV and drawing bloods, the nurse smiles and tells the patient, "OK! Your IV is in and your bloods are all drawn." The EMR documents this time automatically.

Throughout this entire scenario, documentation is done real time, at the bedside, with key information pushed to the treating health care providers. Much of the process is automated, and for those parts which aren't, methods are adapted to minimize substantially the time required to input information into the system. An automated set of checks and balances assure that the physician sees all key information, and further assures that there is an affirmative response to key data, particularly any abnormalities. Once the information is entered, the display is “recalculated” to add this new information to the display in a way such that much information is transmitted via a quick glance at the screen.

4 OUR SYSTEM IN DETAIL

Given the vision presented in the previous section, we now outline the various components of our system that seek to address these ideas in an effective manner. We acknowledge that many commercial as well as research efforts are striving to achieve this functionality in one or the other way. However, these works are still very much evolving, and it is hoped that the ideas we shall present in the following section will add some unique angles to the important overall goal of improving ER healthcare for all.

The overall schematic of our system is illustrated in Figure 1. It consists of the following components: database and knowledge base, inference engine, automatic speech recognizer, information extractor, explanation system, visual processing engine and interactive display system, and authentication/association system. In this paper we focus on the visualization-related system components. Our interface supports multi-touch interaction with the displayed information items (we use an HP multi-touch screen with 1920×1080 resolution accessed with the HP multi-touch API) and voice-controlled information manipulation (using Dragon natural language speech recognition software in the automatic speech recognizer module). The information extractor employs an array of standard and advanced NLP techniques.

4.1 The 5-W scheme of journalistic reporting: the ‘who’, ‘what’, ‘where’, ‘when’, ‘why’, and ‘how’

The framework we have been devising embraces the 5-W principle (who, what, where, when, why, and how) of journalistic

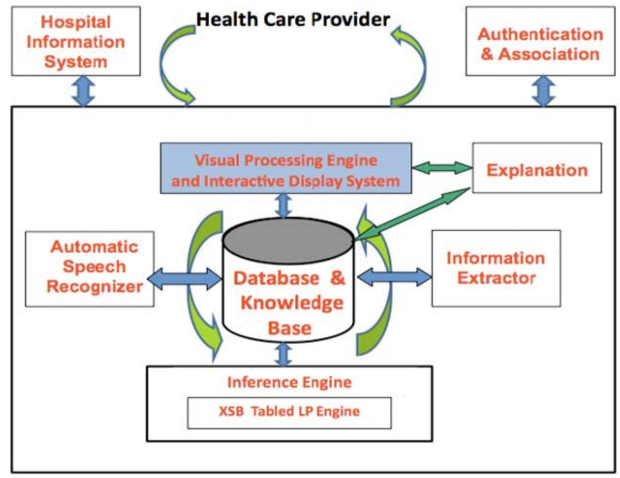

Figure 1: System block diagram.

reporting, covering all important facts and facets on the (medical) case at hand, from dedicated and unique perspectives. By adhering to these principles we aim to aid the physician in the mental categorization and structuring of the complex ER information spaces. More specifically, these are:

- the *‘who’*: the patient (and the doctor, nurse, etc)
- the *‘what’*: the various information categories and items: problems, symptoms, tests and results, diagnosis, treatments, medications, etc
- the *‘where’*: locations (when appropriate) of the ‘what’ on the human body – we use a stylized human body as a map for structured and intuitive information layout
- the *‘when’*: time and duration of the ‘what’ – we use a variation of the well-known LifeLines concept
- the *‘why’*: cause and effect of the various ‘what’ constituents – we visualize this important element of diagnostic reasoning as a directed graph
- the *‘how’*: while somewhat related to the ‘why’ it is also an element of the standardized nurse/doctor-patient dialog/ examination/triage, populating the various information displays

Each of these 5-W uses a dedicated visualization metaphor most suitable to convey its inherent information, such the human map for the ‘what’, the time line for the ‘when’, and the causal graph for the ‘why’. Each of these information displays are linked and can be individually brushed for mutual information selection. We have so far not implemented level-of-detail semantic zooming, but plan to do so in a future implementation. Also, while not described in detail in this paper, our system allows multi-touch interaction with information items on the screen (we use an HP multi-touch screen) and voice-controlled information manipulation (using Dragon natural language speech recognition software). Information extraction from patient text documents uses an array of standard and advanced NLP techniques. In the following we present the various visualization displays of our system within the 5-W decomposition just discussed.

4.2 Visualizing the ‘who’, where’, ‘what’, and ‘how’

These aspects are best described together and in the context of the emergency room clinical encounter outlined in Section 3.2 (with some modifications). Figure 2 (on page 7) illustrates as a sequence of frames the overall user interface of our system as it is populated with information and interacted with during the standardized patient-physician interview. Figure 2a shows the first stage of the clinical encounter: the patient has been triaged by the nurse and some information is already entered into the interface, such as patient name, gender, age, height, weight (the picture is only stylized here). The nurse has found that the patient has elevated temperature (signified in red), but normal blood pressure. This and additional measures, such as complete blood count (CBC) are all organized into a dedicated screen area labelled 'Vitals' on the right. The physician proceeds to the next interview step 'History' (see Figure 2b) and the system signifies this state by highlighting the corresponding interview protocol button (the doctor may always return to this step by clicking the button). The patient describes the reasons for his visit, the symptoms. He seems to suffer under head-related issues. The system highlights the head area on the standardized body map and automatically places the various symptoms as a tag cloud around that area. The next step (Figure 2c) is 'Past History' where the system retrieves patient complaints from previous visits and places them into the interface. Apparently the patient had suffered from a lung cancer and the most recent X-ray is placed next to the lung area on the map, using an illustrated fan to depict the association with this location (its green outline signifies that the cancer is in remission). At that stage, any known allergies are also noted in the corresponding box. The next step in the standardized protocol is the 'Review of Systems' where the doctor asks a series of questions grouped by organ system, such as cardiovascular, respiratory, and so on (nothing of note comes from this step and we skip this frame in Figure 2). Proceeding to 'Social History' (Figure 2d) reveals that the patient has been a smoker (which explains the lung cancer), adding a cigarette icon to the personal information box. After reviewing the 'Family History' which is uneventful, the doctor begins with the diagnosis (see Figure 2e). He first prescribes medication for the fever, which is visually signified by a check mark next to the temperature icon. To diagnose the deeper reasons behind the incidence, he then drags all noteworthy symptoms and data (measurements) into what we call the 'Diagnostic Sandbox' on the top of the screen. This concept has been motivated by the work of Wright et al. [15] and Bier at al. [5] who use the sandbox metaphor as a visual thinking environment supporting both ad-hoc and more formal analytical tasks and sense-making. Likewise we allow the doctor to visually aggregate any displayed data and information items by simple drag and drop operations, and also noting possible diagnoses, here 'epidural hematoma', 'aneurism', and 'head injury'. The doctor then orders the required tests, here a head scan, which is subsequently listed in the 'Pending' box (lower left corner) with urgent or overdue items highlighted in red. Another feature of our interface is the 'News Ticker' located at the bottom of the interface. It forms a centralized and fixed location in which the doctor is alerted to any emerging event, which here could be the patient's rising potassium level or the newly arriving head scan. Finally, Figure 2f shows the head scan fanned to the brain area under examination. The scan clearly shows an epidural hematoma and so the doctor checks off this diagnosis in his hypothesis list in the sandbox.

4.3 Visualizing the 'when'

In the temporal layout, similar to other time-oriented methods the focus is on the time course of the patient's health, letting the physician easily look at the length of a patient's medication, illness etc. This can easily help keep track of test results, and delays if any. It also gives the user a chance to annotate their reasoning by linking events based on causality (see next section).

In the temporal layout visualization, a physician can look at a patient's data from a temporal perspective. The physician can look at all the symptoms reported, the data for the different diagnostic tests, diagnoses and medications. Each *entry* is classified into one of four *tracks* (and assigned the corresponding color):

- Symptoms: Blue
- Tests: Yellow/Orange
- Diagnosis: Green
- Treatments: Pink/Purple

Figure 3 shows the temporal layout for a single patient visit to the physician. Each sub-track within the track represents an event. Both the beginning and the termination of a track have a causal reasoning which is explained by nodes in a panel to the right. These nodes are stacked on the basis of their temporal generation with the bottom most node being the most recent. The nodes are displayed on both sides of the timeline for clarity – the nodes for symptoms and diagnosis on the left, and those for tests and treatments on the right. This minimizes the overlap of the edges connecting the nodes and corresponding events since symptoms and the order of tests tend to happen within a close window. Further, edges which overlap tracks are given a white outline for better trackability.

In the case represented in Figure 3 the patient on admission had a high fever of 103.6F.This symptom starts a fever sub-track within the symptom track. Almost instantly the patient was administered with a dose of Ibuprofen an anti-fever medication simultaneously triggering the generation of another node in the *Treatment* track. Since the drug is administered only once, the event is registered in form of a small square representing an instantaneous event.

A track is drawn until the symptom or medication is reported to be stopped. In this case the fever symptom sub-track is terminated once the patient's fever subsides and is reported to be within normal bounds. In the *Tests* track the start of a sub-track may be triggered by ordering of a specific test and terminated when the results are received. Also, all abnormal results or severe symptoms are highlighted in red in the node panel thus drawing the physician's attention.

A physician may choose to view only one track or more tracks or none at all. To clearly view a path a node can be selected and the path that explains the generation or termination of a track is highlighted. Further, a physican may also choose to associate different subtracks thus forming causal relationships. The physician can later review these associations thus getting a clear picture of the events leading to a diagnosis or treatment. Finally, we note that while we only represent the temporal layout in intervals of 10 minutes, the physician may also zoom out to get a broader temporal layout in hours, days and perhaps years.

4.3.1 A case study

Let us now describe a specific clinical case to show the interactive temporal layout in action (still using Figure 3). Our patient, John Smith, returns to the emergency department. Among other symptoms Smith has now cardiac arrhythmia. This new symptom is laid out within the temporal layout as a sub-track within the symptoms track. A node on the left associated with this sub-track explains the symptom. His physician immediately orders an EKG; this newly ordered test is drawn as a sub-track under the tests track. The physician suspects an ion imbalance and orders a blood chemistry report. This too is drawn under the tests track. Once the physician receives the results for the test, the EKG track is terminated. A node pointing to the end of this terminated sub-track explains the reason for the termination of the sub-track, in this case receipt of the result. The physician receives the result for both the EKG and blood chemistry and these results are

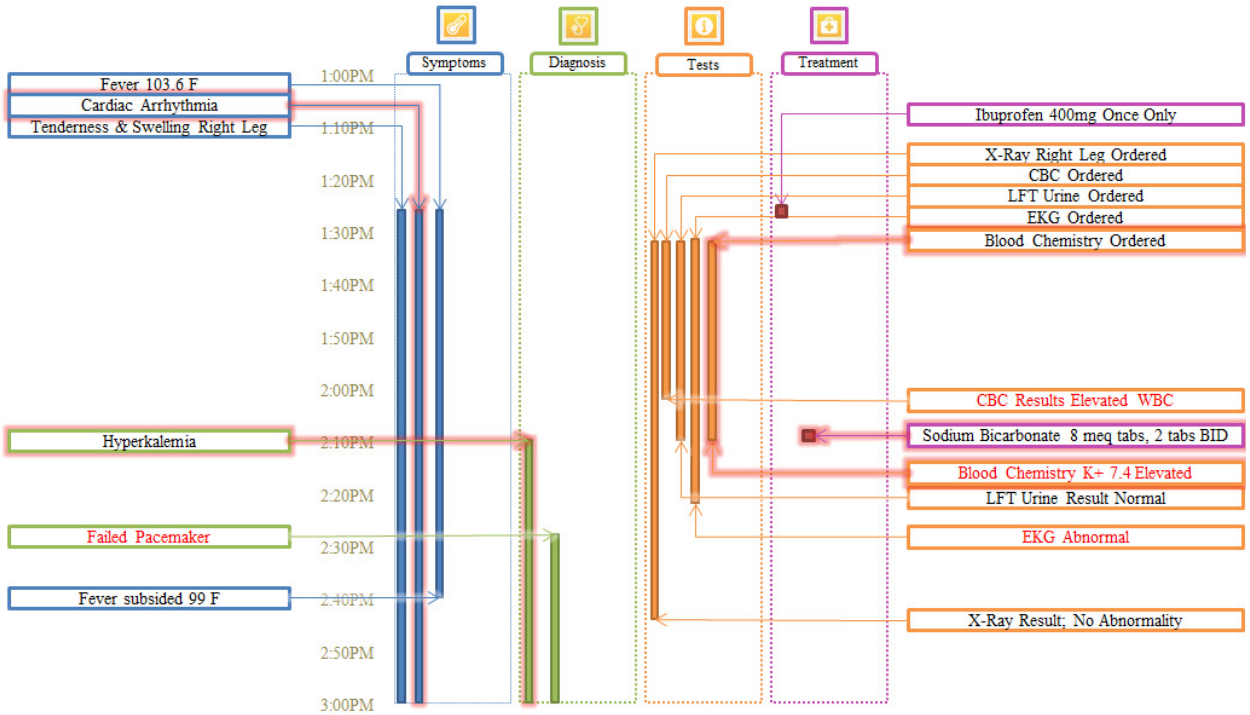

Figure 3: Temporal layout for a single patient visit

highlighted in red to draw the physician's attention to the abnormal results. The physician reviews the test result and diagnoses the patient with Hyperkalemia and also concludes that the patient's pacemaker has malfunctioned. These two diagnoses are laid out as sub-tracks within the diagnosis track. The patient is placed on medication as an intervention for Hyperkalemia, which now appears under the treatment track. Now the physician can start marking causal relationships between the nodes in the layout. The causal path leading to the diagnosis of Hyperkalemia and its related medication is highlighted in Figure 5. As time progresses his potassium levels do not seem to normalize which causes the sub-track to keep running and so does the diagnosis sub track for the failed pacemaker. These sub-tracks run until the condition has been taken care of and are no longer a reason for concern. These sub-tracks constantly remind the physician of the conditions a patient is suffering from. A physician may also associate different events such as Hyperkalemia and the drug intervention associated with it forming causal relationships. This gives the physician the ability to later review his decisions in a temporal layout.

4.4 Visualizing the 'why'

The representation of the why, that is, the logical chain of reasoning of the physician can be done compactly using a directed graph layout. Such a layout can help detect errors in diagnoses, as well as help a new physician come up to date with a patient's symptoms and illnesses very quickly.

In the causal graph layout, the physician can get a quick overview of the causality relations between the different sub-tracks presented in the temporal view. Here, each sub-track is represented as a node, and there is a *directed* edge between any two nodes n_1 and n_2 if n_1 causes n_2. For example, in Figure 3, we see the layout for a single patient visit, with the date at the center. Here each type of node (sub-track) is assigned a color similar to the one in the temporal layout. Further, to link the symptoms together, they have incoming edges from a node representing a visit to the physician. Such 'incident' nodes are red in color, and show the date of the visit. The graph layout is done using a modified version of the force-directed layout [7]. We have used the Prefuse Visualization Toolkit [9] to create all of our layouts. Force-directed layouts focus on positioning the nodes of a graph so that all the edges are of more or less equal length and there are as few crossing edges as possible. The edges between nodes in different cases are allowed to be longer for a better layout. Further, we draw a convex hull around the nodes belonging to a single visit, showing a clearer delineation between different visits.

Figure 4 shows the causality graph for a single visit that happens on 6/10/2007 (the same as the temporal graph in Figure 3). A generic causality path follows the following format – a patient has multiple symptoms for a visit. One or more symptoms lead to test(s), and finally the test results lead to a diagnosis followed by a treatment. In order to avoid multiple nodes to represent a sub-track, we use visual cues to represent the current state of the track. The nodes are drawn fainter if the treatment or diagnosis is not valid any more for the patient, or if the test results are pending (see the node for X-Ray in Figure 4). Further, for test nodes, its boundary is drawn depending on the severity of the result – none for normal, blue for mild severity, and red for high severity. To get a quick overview of the sub-track (as represented in the nodes in the temporal graph), the physician can hover over it to see the details related to it. Further, in order to link current symptoms to past diseases and medications, the physician can ask for the list of side effects as well.

4.4.1 A case study

Let us return to our fictitious patient John Smith. In Figure 5 (see page 8), the physician is looking at the multiple symptoms of Mr. Smith who comes into the emergency room on 10/21/2010. To get an idea of the patient history, the physician pulls up cases in the recent past to look at the current illnesses and treatments the patient might be under. At this stage the treatments prescribed during the visit in 2007 (shown in Figure 3 and 4) are no longer being used – the pacemaker was fixed, and the patient was cured of Hyperkalemia. However, as the physician looks at the visit on 8/16/2009, he notices that Mr. Smith was diagnosed with Multiple Sclerosis, Diabetes, and Hypothyroidism. These being long-term diseases could possibly explain some or all of the symptoms the patient is suffering from now. Further, the medications which follow these diagnoses can also cause side effects leading to new symptoms.

As the physician considers the symptoms one by one (see Figure 5a), the presence of fever and "inflammation and redness" lead to the diagnosis of Cellulitis, and the patient is put on an antibiotic. Further, the physician also links the Cellulitis to the patient's precondition Diabetes via a causal link.

Next as he looks at the symptom for joint pain (the only symptom unresolved in Figure 5a) he hypothesizes that it might be related to a medication, Avonex that was previously prescribed (see the 8/16/2009 cloud). So the physician pulls up a listing of the side effects for this drug (the yellow text boxes in Figure 5b). He notices that among the new complaints, the joint pain could indeed be caused by Avonex. An alternative hypothesis might be that the joint pain is due to bone weakness, and to make sure the physician calls for a bone density test (BMD) and an X-Ray. The combined results indeed show a low bone density leading the physician to pronounce the diagnosis of Osteoporosis, and medication for the same (see Figure 5c).

Other issues are the symptoms of coughing and difficulty of breathing which make the physician suspect trouble with the thyroid gland. He orders a test for the thyroid gland (TSH/T4). The coughing also leads the physician to order an ultrasound. The test results point to Goiter. Further, the results for TSH/T4 are linked to the preexisting condition of Hypothyroidism, thus serving as an explanation for the abnormal test values.

Finally, the multiple symptoms of frequent urination, dysuria and lower abdominal pain suggest a Urinary tract infection to the physician. He orders a Urinalysis test whose results confirm his suspicion. In addition, one of the symptoms in this set (frequent urination) is linked to the patient's diabetes.

This example clearly demonstrates that multiple symptoms along with knowledge about the history can lead to a quicker diagnosis, and can assist the physician in ordering only the most relevant tests. Further, the links between past medication/diagnosis, and current symptoms also goes some way into explaining "how" the patient got the new set of symptoms in the first place. Finally, incorporating (medication) information from a backend database is also very useful to the physician. We believe that such a system can be further extended to store logical relations from current physician diagnoses and other expert sources.

5 CONCLUSIONS

We have a described a fledgling visual analytics system for deployment in an emergency department setting. We aimed for a one-stop information interaction framework where the physician is fully in charge, making informed decisions based on his own expertise and intuition. All components have been developed in collaboration with emergency physicians at a large teaching hospital, and our next steps will include more formal user and affordance studies to fine-tune the various modules of our system. We are currently implementing a tighter integration of these in terms of look and feel. Further, we are working on a causal graph layout that is more geometrically structured than the present free-form force-directed layout, and on an improved temporal plot.

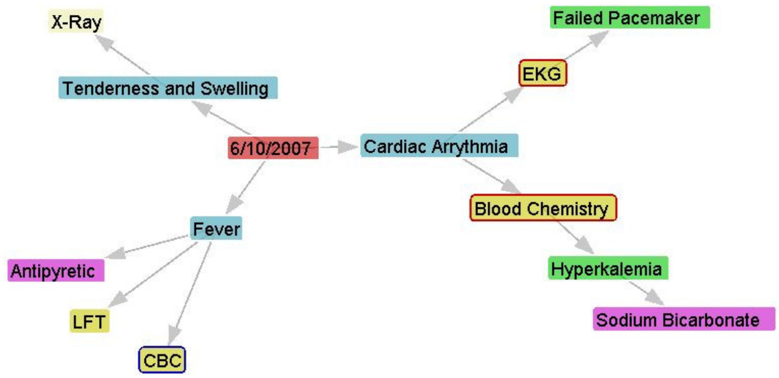

Figure 4: Causal graph layout of the patient visit as shown in Figure 2

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. van Bemmel and M. Musen, ed., Handbook of Medical Informatics, Springer, The Netherlands, p. 99, 1997.

[5] E. Bier, S. Card, J. Bodnar, "Entity-based collaboration tools for intelligence analysis," VAST Symposium, pp. 99 – 106, 2008.

[6] 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.

[7] T. Fruchtermann, E. Reingold, "Graph drawing by force-directed placement," *Software – Practice & Experiments*, 21(11):1129-1164, 1991.

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

[9] J. Heer, S. Card, J. Landay, "Prefuse: A Toolkit for Interactive Information Visualization," *Proc. ACM Human Factors in Computing Systems (CHI)*, pp. 421-430, Apr 2005 (toolkit available at http://prefuse.org)

[10] 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.

[11] 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.

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

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

[14] 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

[15] W. Wright, D. Schroh, P. Proulx, A. Skarbuskis, B. Cort, "The Sandbox for Analysis: Concepts and Methods*." Conference on Human Factors in Computing Systems*, pp. 801-810, 2006.

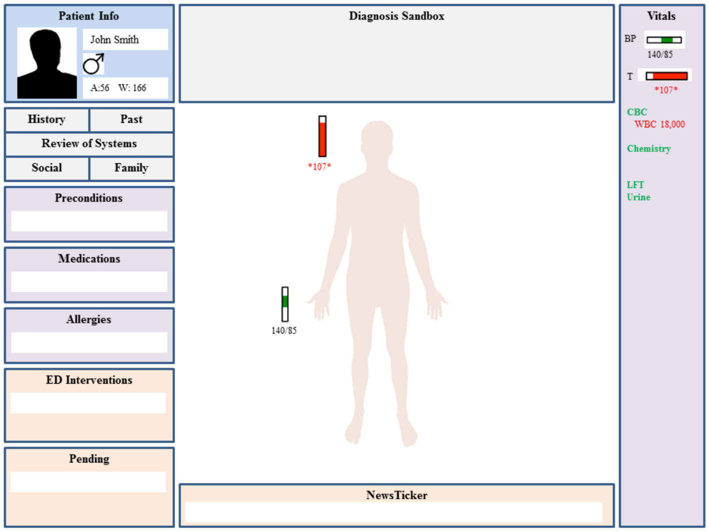

(a) Results of triage (name, weight, age, vitals with high fever)

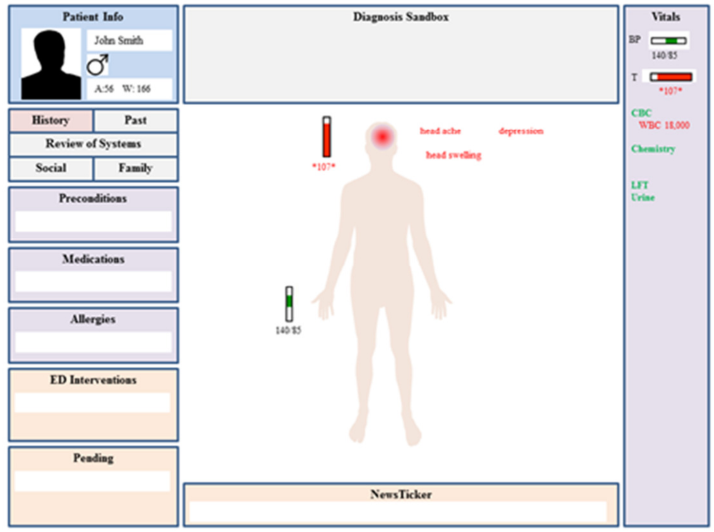

(b) Populating the spatial map with symptoms (head problems)

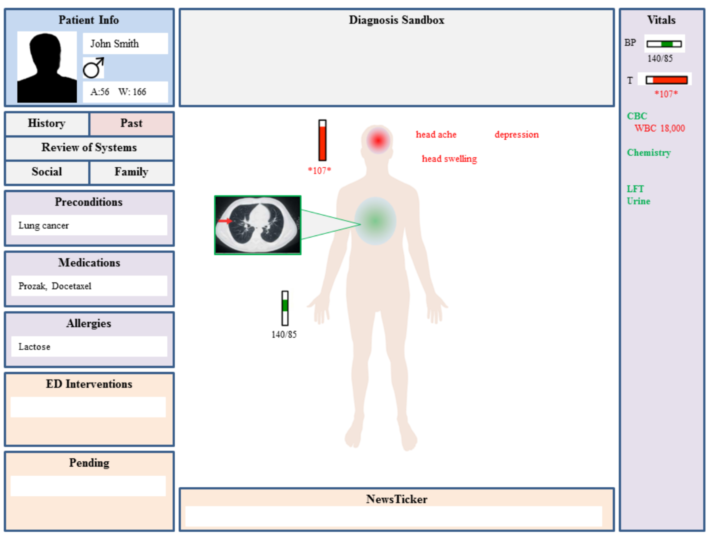

(c) Adding past history (lung cancer in remission)

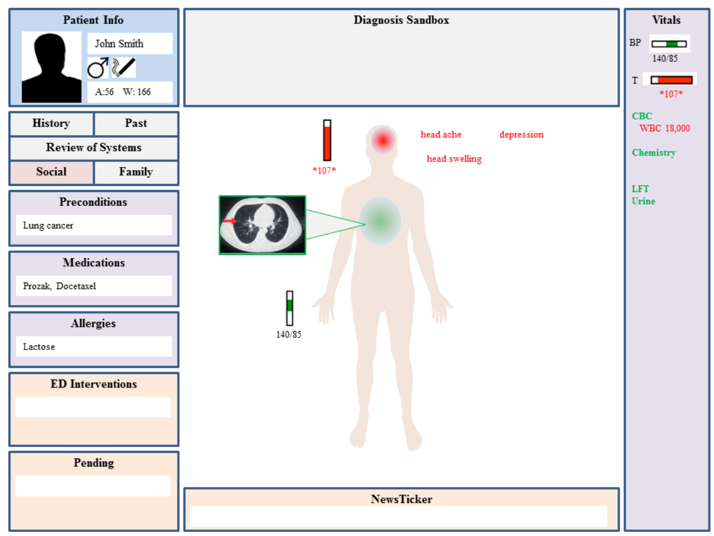

(d) Social history indicates patient has a history of smoking

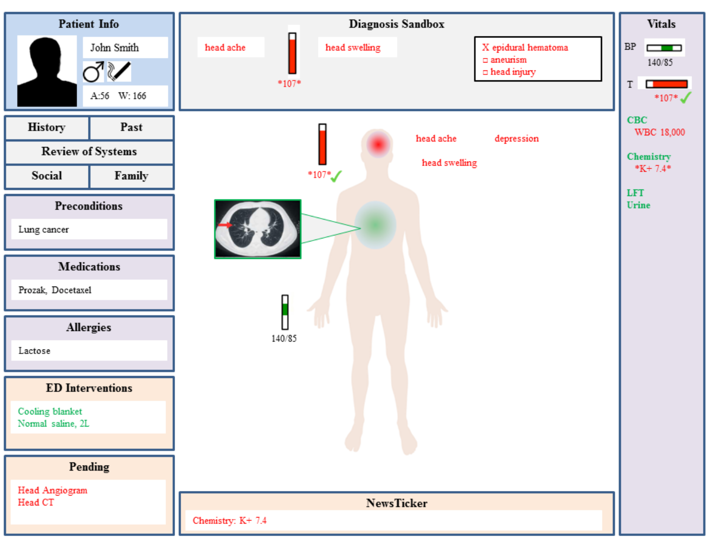

(e) The diagnostic process begins, using the Diagnostic Sandbox

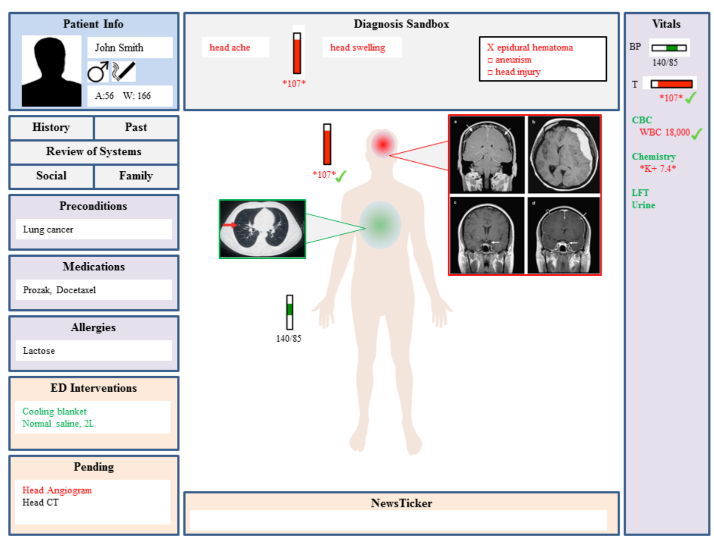

(f) The diagnostics determines epidural hematoma

Figure 2: Overall user interface as it is populated with information and interacted with during a patient-doctor interview.

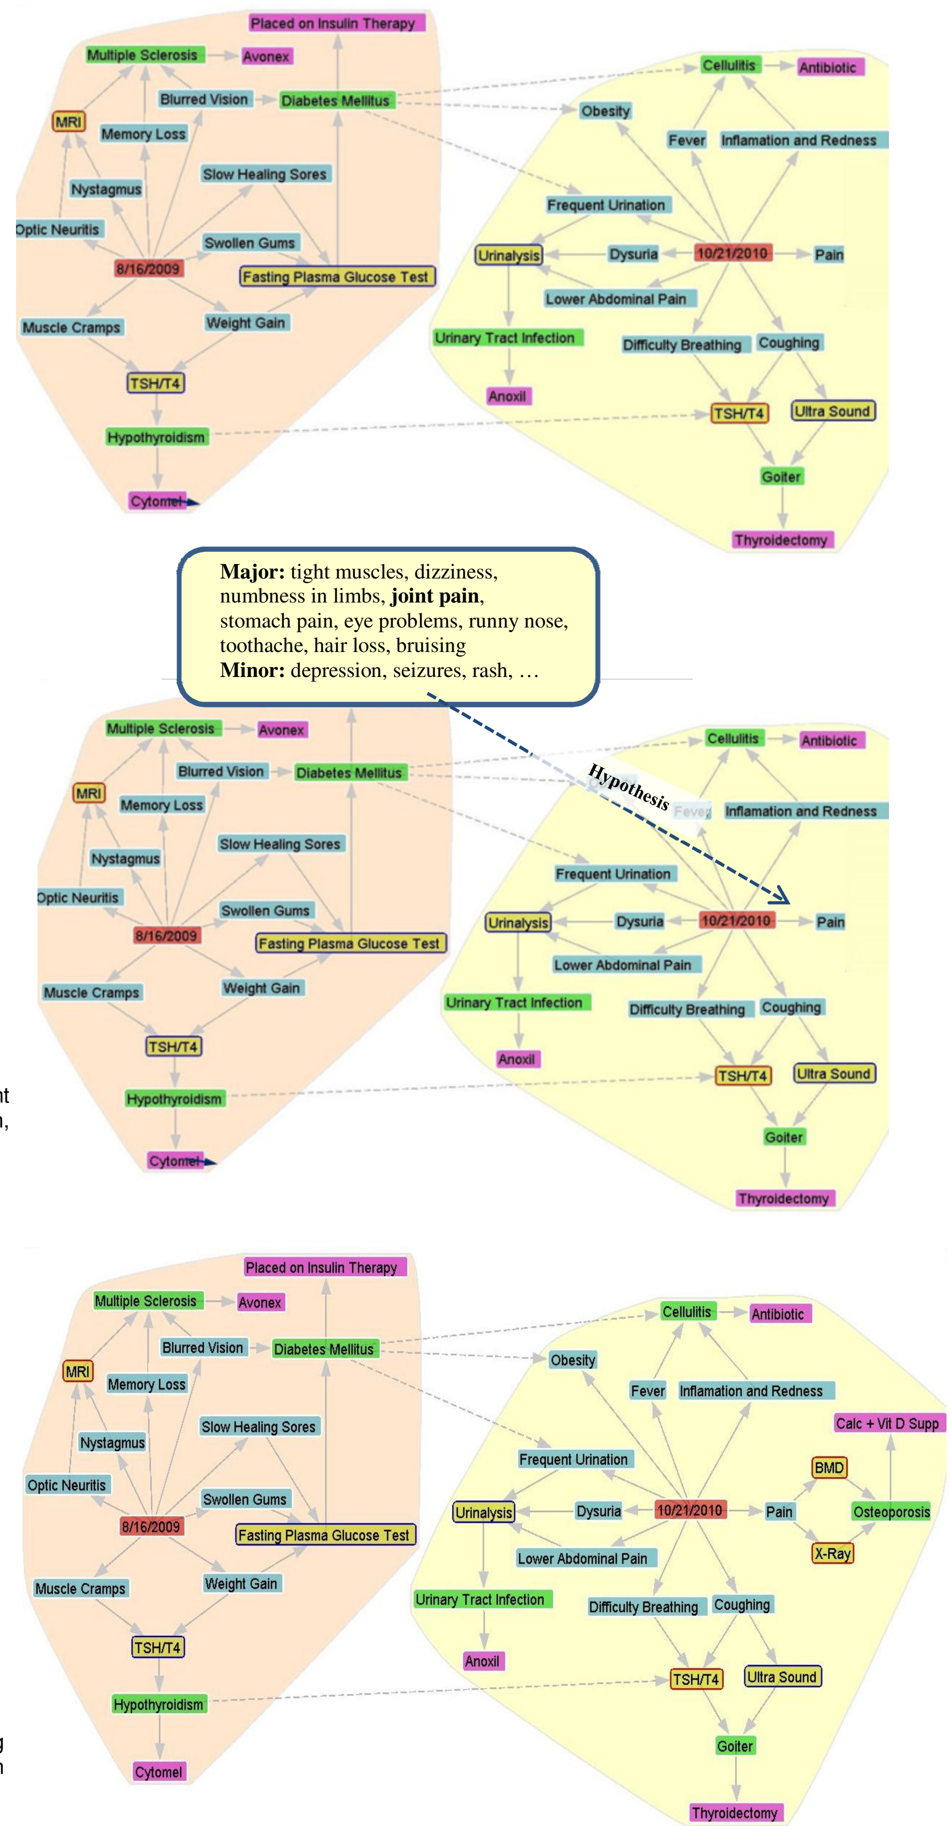

(a) The physician has diagnosed and treated all symptoms except for 'joint pain'

(b) The physician is drawing a hypothesis for joint pain based on the side effects of a medication, Avonex.

(c) The physician arrives at a conclusion based on the test results

Figure 5: An example showing the reasoning process for multiple symptoms presented on 10/21/2010.