

Context-Based Electronic Health Record: Toward Patient Specific Healthcare

William Hsu, *Member, IEEE*, Ricky K. Taira, Suzie El-Saden, Hooshang Kangarloo,
and Alex A. T. Bui, *Member, IEEE*

Abstract—Due to the increasingly data-intensive clinical environment, physicians now have unprecedented access to detailed clinical information from a multitude of sources. However, applying this information to guide medical decisions for a specific patient case remains challenging. One issue is related to presenting information to the practitioner: displaying a large (irrelevant) amount of information often leads to information overload. Next-generation interfaces for the electronic health record (EHR) should not only make patient data easily searchable and accessible, but also synthesize fragments of evidence documented in the entire record to understand the etiology of a disease and its clinical manifestation in individual patients. In this paper, we describe our efforts toward creating a context-based EHR, which employs biomedical ontologies and (graphical) disease models as sources of domain knowledge to identify relevant parts of the record to display. We hypothesize that knowledge (e.g., variables, relationships) from these sources can be used to standardize, annotate, and contextualize information from the patient record, improving access to relevant parts of the record and informing medical decision making. To achieve this goal, we describe a framework that aggregates and extracts findings and attributes from free-text clinical reports, maps findings to concepts in available knowledge sources, and generates a tailored presentation of the record based on the information needs of the user. We have implemented this framework in a system called *Adaptive EHR*, demonstrating its capabilities to present and synthesize information from neurooncology patients. This paper highlights the challenges and potential applications of leveraging disease models to improve the access, integration, and interpretation of clinical patient data.

Index Terms—Data visualization, health information management, knowledge representation, natural language processing (NLP).

I. INTRODUCTION

ELECTRONIC health records (EHRs) provide a centralized location for aggregating patient data acquired from

different sources and at multiple biological scales, with the aim of making this data readily accessible to healthcare professionals. While an intent of digitizing health records has been to lower the cost of healthcare, reduce the number of preventable medical errors, and improve the accuracy of diagnosing and treating patients [1], the sheer amount of data collected poses new challenges. Physicians often need to strike a balance between managing a large number of patient cases and spending sufficient time to thoroughly review a patient's medical history. One study showed that the volume of work associated with primary care visits has increased, resulting in a shorter amount of time available to address individual tasks such as diagnosing patients, prescribing medications, ordering procedures, and providing counseling or physical therapy [2]. Today, a comprehensive review of the patient's health record would require a clinician to examine documents, medical images, and charts while mentally noting issues related to the current clinical context—all while disregarding unrelated information contained within the presented electronic record. Given their time constraints, clinicians are limited in their abilities to process all of this data simultaneously [3]. As such, much of their time is spent skimming parts of the patient record until useful information is found [4]. This problem is compounded by the addition of new information derived from genomic analyses, which provide additional evidence that needs to be understood and interpreted in the context of the entire patient record.

This paper addresses the challenge of integrating fragments of information captured in clinical reports, laboratory tests, and ordered procedures. In current EHR implementations, results and interpretations are often scattered across different parts of the user interface, requiring a user to navigate through multiple screens to find relevant information. For example, if a neurooncologist wishes to determine if her patient is eligible for a clinical trial, she would need to review multiple documents such as oncology reports (family/treatment history/outcome measures), radiology reports (evidence of tumor progression), pathology reports (grade, histological type), and laboratory results (e.g., serum creatinine levels). As EHRs become the primary repository for all clinical data generated, adapting the presentation of this data to different users becomes more important [5]. Unique views of patient data can be created to meet the information needs of each user. A concise presentation that integrates relevant information across all available sources and displays only the necessary details would not only help practitioners reduce the time spent searching for relevant data, but also assist them to utilize this data more effectively to inform personalized care and medical decision making.

Manuscript received April 15, 2011; revised October 31, 2011; accepted January 16, 2012. Date of publication February 3, 2012; date of current version March 7, 2012. This work was supported in part by the National Institute of Biomedical Imaging and Bioengineering under Grant R01-EB009306 and by the National Library of Medicine under Grant R01-LM009961.

W. Hsu is with the Department of Radiological Sciences, University of California, Los Angeles, CA 90024 USA (e-mail: willhsu@mii.ucla.edu).

R. K. Taira, H. Kangarloo, and A. A. T. Bui are with the Department of Radiological Sciences, University of California, Los Angeles, CA 90024 USA (e-mail: rtaira@mii.ucla.edu; kangarloo@mii.ucla.edu; buia@mii.ucla.edu).

S. El-Saden is with the University of California, Los Angeles, CA 90024 USA, and also with the Department of Radiology, Greater Los Angeles Veterans Administration, Los Angeles, CA 90073 USA. (e-mail: sels@mednet.ucla.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TITB.2012.2186149

We present an application called Adaptive EHR (*AdaptEHR*), which is built upon a context-based framework, integrating data from different sources in the patient record and subsequently tailoring the presentation based on the contents of the patient's clinical reports and the information needs of the user. This paper reviews past developments on integrating and displaying multimedia patient records. Next, a context-based framework is presented that utilizes two sources of knowledge: biomedical ontologies and graphical disease models. Lastly, *AdaptEHR* is presented and discussed in the context of neuro-oncology.

II. BACKGROUND

The way a patient's health record is presented can affect a clinician's ability to understand and interpret the information. Poor integration between the various data types (e.g., images, text, numerical data) and the inclusion of excess information can hamper physicians in noticing important findings. Determining the optimal way to organize and present the patient record has been a widely researched area with many efforts focused on providing a summary of the patient's information using compact and intuitive graphical representations. Initial work focused on the addition of images, audio, and charts to text-based displays [6], [7]. Powsner and Tufte presented a method for summarizing the patient record using a series of graphs on a single page [8]. While the summary effectively depicted trends in the data, it lacked interactivity and only supported numerical data that could be graphed. Plaisant *et al.* developed an interface for plotting events, images, and values on a timeline [9]; more recently, the timeline representation has been augmented to allow users to align, rank, and filter data based on meaningful events in the record [10]. Finally, Bui *et al.* developed an interactive TimeLine visualization that filters data elements by medical problem. The novelty of their work is its ability to integrate multiple parts of the patient record based on a medical problem and display only portions of the record that are related to that problem [11]. *AdaptEHR* can be considered an evolution of TimeLine. It is similar to TimeLine in the way it organizes the health record around problems and findings and utilizes timelines to present events in the patient record. However, while TimeLine requires manually coded rules to determine how data are displayed, *AdaptEHR* attempts to automatically infer these rules and relationships through the use of knowledge from biomedical ontologies and probabilistic graphical models.

One early example of an application that utilizes domain knowledge to generate tailored views of medical data is described in [12]. It incorporates an ontology called the medical entities dictionary (MED) [12], which represents common terms from four hospital systems (laboratory, electrocardiography, health records coding, and pharmacy). For a given concept (e.g., heart), the system utilizes MED to retrieve other related data from the health record (e.g., references to heart diseases, cardiac enzyme tests, chest X-ray). In contrast, rather than relying on a knowledge source, the NeuroAnalyzer system utilizes a machine learning approach to automatically cluster clinical reports around specific topics (e.g., documents related to the brain) [13]. A recent survey of current commercially and in-

house developed EHRs provides a variety of functions to support data integration and decision support [14]. Many systems provide a summary view of a patient's record, but users can only select a single tab or link at a time to view available data types (e.g., chart review, laboratory results, clinical notes). While systems have incorporated methods for data standardization using controlled vocabularies, functionality to utilize ontologies to annotate patient data is not currently available. Finally, links to scientific literature are also being provided through tailored recommendations using systems such as InfoButtons [15], but such an approach only provides a link to the resource and does not attempt to provide an interpretation of how findings from the study can be personalized for a specific individual's case.

III. SOURCES OF DOMAIN KNOWLEDGE

Arguably, a disconnect exists between the acquisition of knowledge (i.e., the testing and validation of a clinical hypothesis) and the application of this knowledge (i.e., evidence-based medical practice). Incorporating knowledge sources into the EHR to provide context would be a step toward bridging this gap. For example, physicians can leverage evidence from experimental studies to identify imaging-based biomarkers that predict treatment efficacy and improved outcomes. In addition, linkages among clinical observables, physiological processes, signaling pathways, and cellular/genetic factors may be established and used to generate explanations for specific patient presentations. *AdaptEHR* leverages two types of knowledge sources.

A. Biomedical Ontologies

Ontologies capture knowledge about the medical domain by representing relevant concepts and their relationships in a structured, reusable, machine-understandable format. We incorporate ontologies to address three aims.

- 1) *Data integration*: The vocabulary that is used to describe findings is highly variable and dependent on the physician. While such flexibility allows a rich set of terms to be used to communicate findings, it poses a challenge when disambiguating whether multiple utterances of a finding reference the same instance. Ontologies provide a means for mapping synonymous terms to unique concepts.
- 2) *Information retrieval*: An ontology also supports query expansion, which augments the original query with synonymous or related terms to improve recall of items that are relevant but may not be literal matches to the original query [30]. For example, if a physician wishes to identify all findings related to the temporal lobe of the brain, an ontology can be used to identify other related structures that are part of the temporal lobe such as the hippocampus, collateral eminence, and amygdala.
- 3) *Data annotation*: Finally, ontologies contain information that can be used to annotate clinical data. For example, a mention of epidermal growth factor mutation (EGFRvIII) extracted from a pathology report can be mapped to the standardized term "EGFRvIII peptide" and annotated with a semantic type derived from the Unified Medical Language System (i.e., amino acid, peptide, or protein), a

definition from the National Cancer Institute Thesaurus (NCIt) (i.e., a synthetic peptide sequence. . . caused by the deletion of exons 2-7), and associated with a parent concept (i.e., vaccine antigen).

While ontologies attempt to provide sufficient coverage within a defined domain, to date no single ontology provides sufficient coverage for standardizing and annotating the entire patient record. Given our application domain, we have selected ontologies that provide broad coverage in the areas of anatomical location (Foundational Model of Anatomy), image characteristics (RadLex), clinical terms (Systematized Nomenclature of Medicine—Clinical Terms), cancer-related findings (NCIt), and medications (RxNorm). Phrases extracted from clinical reports are mapped to ontologies using tools such as the National Center for Biomedical Ontology’s BioPortal annotator Web service, which encodes results in eXtensible Markup Language (XML) [16].

B. Graphical Disease Model

While biomedical ontologies provide breadth in coverage, a model that captures detailed characteristics of a disease is necessary. The purpose of such a model would be to: 1) enumerate all of the common types of findings, tests, and procedures that are reported for a given disease; 2) provide (causal) explanation of how represented elements relate to one another; and 3) permit computational analysis of the influence among variables based on encoded probabilities. Graphical models—Bayesian belief networks (BBNs), in particular—are powerful because they combine both graph and probability theories to encode qualitative and quantitative knowledge [17]. Intuitive to interpret, a BBN’s graph structure also encodes conditional independence relationships among variables. The probabilities not only specify the degree of dependence among variables, but also enable a class of efficient algorithms to perform reasoning on the model [18]. We have previously reported efforts toward building a Bayesian network consisting of clinical and imaging features to predict the outcome of brain tumor patients and osteoarthritis using a BBN [19], [20]. We briefly examine three properties of the BBN that are used to provide context.

- 1) *Variables*: BBNs contain two types of variables: target and evidence variables. Target variables represent outcomes that are of particular interest to the user (e.g., survival). Evidence variables represent various observations or intermediate stages that have an effect on the target variable. They can be identified from a variety of data types (e.g., oncology reports, imaging, pathology) and biological scales (e.g., genetic, cellular, organ). Each variable is associated with a set of discrete states. For example, our model includes a variable for *necrosis*, which contains two states: *present* or *absent*. Variables are also derived from findings reported in clinical reports. We can also add variables to the BBN corresponding to previously unseen findings found in clinical reports or literature.
- 2) *Relationships*: The structure of the model specifies how variables relate to one another. In a BBN, relationships are represented as a directed edge connecting two vari-

ables. The graph structure conveys how information flows across variables in the network. The *Markov assumption* states that a variable is independent of its ancestors, given knowledge about its parents; this property leads to a rule called the *Markov blanket*, which identifies the set of variables that need to be specified to fully characterize a target variable. This property can be used to identify a set of findings related to the target variable of interest that should be presented together in the display.

- 3) *Parameters*: Associated with each variable is a set of probabilities that is stored in a conditional probability table (CPT). The CPT characterizes the probability of a variable being a certain state; probabilities for a variable may change based on the known states of other variables and the structure of the model. Using the CPTs, we can quantify how strongly a target variable is affected by another variable by removing the connection between the two variables and comparing the resulting joint probability distributions [21]. The comparison can be made using the Kullback–Liebler divergence [22], and the resulting value is called the *strength of influence*: a greater divergence value reflects stronger influence between the two variables, providing a quantitative means for identifying findings that are closely related. Based on the selected target variable, findings are ordered based on the influence measure.

A subset of a neurooncology model is depicted in Fig. 1. An important source of domain knowledge that is translated into the variables, relationships, and parameters of the model comes from the results of controlled trials and experimental studies. We are developing methods to incorporate results from published research studies [23]. We use keywords such as “immunohistochemistry AND (glioblastoma multiforme OR GBM) AND (prognosis OR prognostic)” to search PubMed and obtain an initial set of publications that are then filtered to include only clinical trial papers published in the past 10 years (i.e., 2001 to 2011). Papers are assessed using criteria developed by the Grading of Recommendations Assessment, Development and Evaluation working group, which rates the quality of evidence presented in a paper by criteria such as design, quality, consistency, and directness [24]. Each study is manually inspected, extracting information such as:

- 1) study hypothesis and observed (outcome) variables;
- 2) study population characteristics;
- 3) experimental method, including specifics of assay techniques, specific platforms, normalization methods;
- 4) statistical tests;
- 5) study conclusion, including hypothesized pathways and explanations relating the study variables and outcomes.

The extracted information is used to augment the model with additional variables and provide annotations (e.g., study population, statistical method) to relationships specified in the model that can potentially be used to generate explanations. Presently, our model incorporates variables representing clinical observables, image findings, pathology, treatments, and genomic analysis. The model is in the form of a concept graph (the CPTs have not been computed yet), and we are exploring methods to

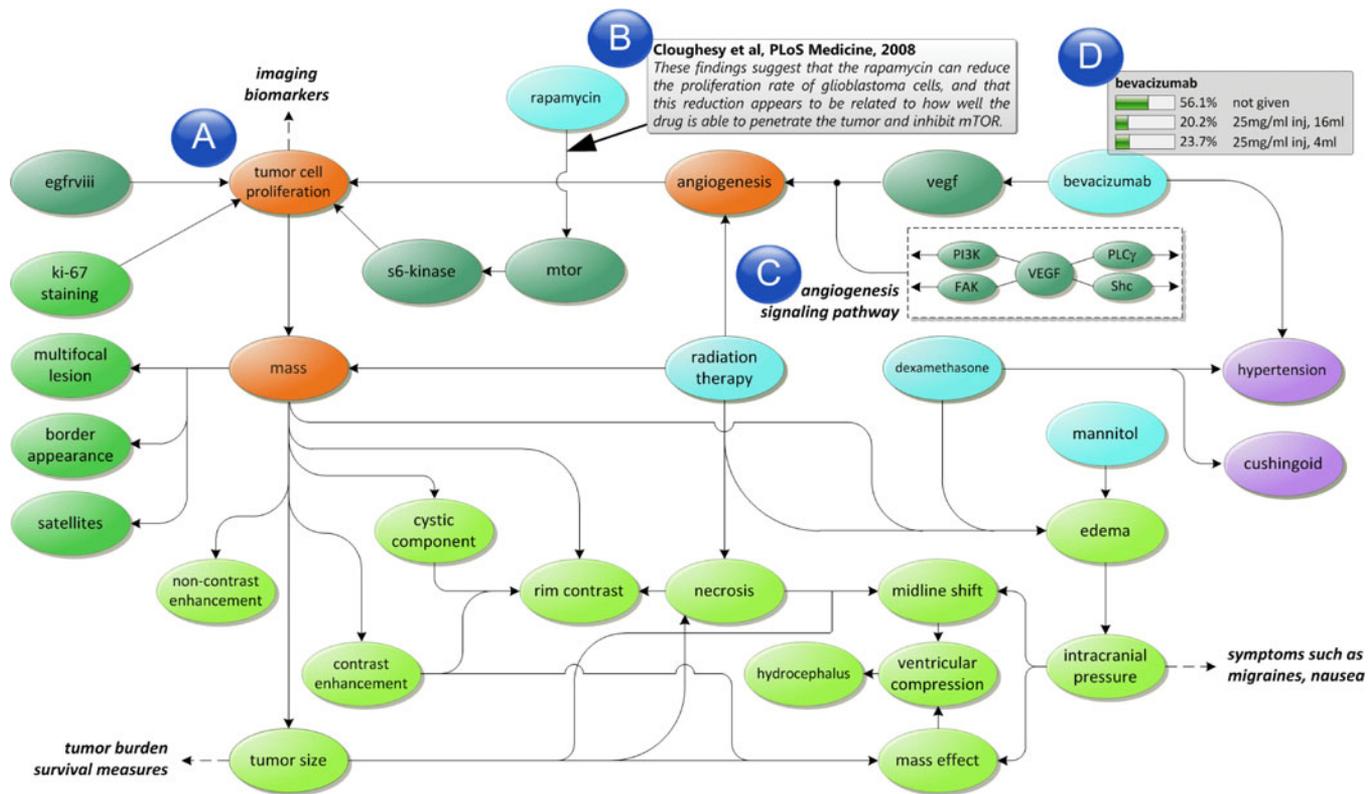


Fig. 1. Example of a graphical disease model for neurooncology. (a) Model incorporates variables from different sources characterizing multiple biological scales. (b) Each relationship is annotated with references to scientific literature. (c) Additional contextual information is provided by incorporating signaling pathways that are used to generate explanations. (d) Each variable is associated with a table that enumerates the possible states and conditional probability distribution for that variable given its parents.

estimate probabilities from structured patient data (using algorithms such as expectation maximization) or from values reported in literature using statistical meta-analysis.

IV. CONTEXT-BASED FRAMEWORK

The process of generating a context-based view of the patient’s health record can be summarized in three steps, as illustrated in Fig. 2: 1) aggregate the patient data from separate clinical sources; 2) structure the patient record to identify problems, findings, and attributes reported in clinical reports, and map them to available knowledge sources; and 3) generate a tailored display based on annotations provided by the knowledge sources.

A. Aggregating Patient Data

The health records for 283 neurooncology patients with a confirmed diagnosis of GBM, an aggressive form of brain cancer, were electronically retrieved from hospital information systems using a system called *DataServer* [25], which is a distributed infrastructure for querying multiple clinical data sources and generating a uniform representation of their outputs. Institutional review board approval was obtained prior to the collection and processing of the data. Only radiology reports (with images), pathology, oncology notes, consultation letters, surgical notes, admission/discharge summaries, demographics, and laboratory

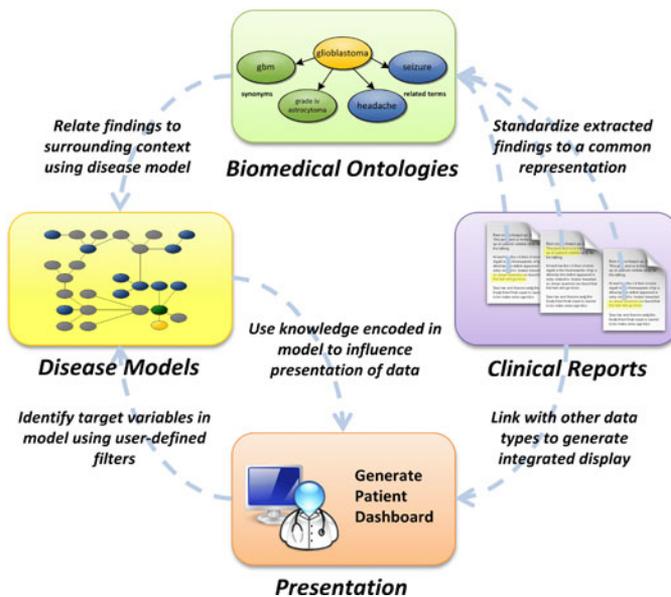


Fig. 2. Flow of information in the context-based framework.

results were examined. The entire dataset was used to identify new findings, guide variable selection, estimate values for the probabilistic disease model, and provide a test set for validating the system.

B. Structuring the Patient Record

For each patient record, we follow a methodical approach to extract problems, findings, and attributes from clinical narratives, mapping them to concepts in the knowledge sources. While a natural language processing (NLP) system [26], [27] that has been trained to extract biomedical concepts (e.g., disease, finding) and attributes (e.g., duration, existence, location) from semistructured medical documents is used to automate some tasks, the abstraction task is primarily dependent on a human annotator to oversee the process.

- 1) *Identifying findings*: All problems/findings and references to anatomical locations in each clinical report are highlighted using NLP, which results in a list of finding utterances presented to the user. For example, in the sentence, “There is an enhancing mass in the left parietal lobe,” the finding “mass” and anatomical location “left parietal lobe” would be automatically identified by the NLP system. The user would then confirm that “mass” is a finding, confirm its location (i.e., left parietal lobe), and specify the spatial relationship characterizing the location description (i.e., in).
- 2) *Characterizing attributes*: In addition, for each finding, the system presents all linked sentences and a specific object frame that allows precise characterization of the finding’s attributes. Each finding is presented as a frame, and its slots represent its attributes. Each frame and slot is dictated by variables, states, and annotations represented in the disease model.
- 3) *Performing coreference resolution*: For each confirmed finding, the user specifies if an utterance is the first mention of the finding or if it is associated with a previously mentioned instance of that finding. The system suggests a coreference based on matching anatomical locations. Once findings have been identified and structured from all reports, the user is asked to correlate descriptions of findings across reports so that each individual finding can be represented with respect to time. The system presents a list of prior findings of a similar type (from previously structured/reviewed reports), allowing the user to create and validate links. The result of this process will be a set of linked, formalized representations of each finding represented over time.

The resulting structured output is encoded in an XML file. Thus far, 20 patient cases have been structured representing a total of 717 documents. The initial case took two days to structure approximately 200 reports due to time spent adding new variables to the model. Further details on the structuring tool and its ability to identify findings are reported in [28].

C. Interpreting Patient Data Using Available Knowledge

Once findings and attributes have been extracted from clinical reports and mapped to variables in the graphical model, the next step is to leverage the model structure and parameters to identify relationships and generate an integrated display. In this paper, we explore two ways to utilize the domain knowledge.

- 1) *Integrating across data types*: The traditional way that EHR user interfaces present patient data is to either organize them by source (e.g., radiology, pathology) or time (e.g., when the report was generated). In these approaches, physicians are often left to synthesize information found in different parts of the record. The context-based approach leverages the disease model to link related data from any sources and time periods to the same variable in the model. For example, the model can be used to link image findings (e.g., rim enhancement) found in the raw image data (e.g., segmented regions) and radiology reports (e.g., mentions of enhancement).
- 2) *Identifying associations between findings*: Different data sources provide multiple perspectives of the same finding. For example, swelling of the brain (edema) can be characterized by imaging characteristics (e.g., appearance of mass effect and/or midline shift), neurological exam (e.g., whether the swelling affects motor skills), and prescribed medications (e.g., dexamethasone, which is used to combat the swelling). Once findings from the patient record are mapped to concepts in the disease model, path analysis techniques can be used to characterize how findings are related based on the model’s topology.

V. ADAPTEHR APPLICATION

We have implemented a functional prototype of the context-based architecture in an application called AdaptEHR, which is written in Java using Swing components and is connected to a relational database (i.e., MySQL) that contains a structured and annotated representation of each patient’s record. While we present the functionality of AdaptEHR in the context of neurooncology, it can be generalized to any disease given the requisite knowledge sources. The interface, as depicted in Fig. 3, consists of the following components.

- 1) *Patient dashboard*: The patient dashboard comprises multiple panes that summarize the contents of the patient record. The dashboard presents information such as past encounters, prescribed medications, findings documented in clinical reports, and results of laboratory tests and procedures. Clinical findings, which have been extracted by the NLP-assisted structuring tool, are presented along a timeline, which visually summarizes when each finding is mentioned in a report. Each observation of a finding is also color coded based on an interpretation of whether the finding has improved, worsened, remained unchanged, or not present [29].
- 2) *Query pane*: The query pane provides a set of controls that allow users to specify target variables (e.g., to determine whether a chemotherapy regimen is effective), define a temporal window (e.g., show only observations within a six month period after the patient’s surgical resection), group findings by semantic group (e.g., findings of the nervous system), and filter by source (e.g., show only radiology reports) or type (e.g., display only magnetic resonance images).



Fig. 3. Screen capture of the AdaptEHR patient dashboard. (a) Findings list serves as the primary method of navigating the patient record. (b) Timeline summarizes each observation of a finding across time, color coded by trends such as improving (green), worsening (red), present (gray), and not present (white). See online version for color figure. (c) Findings are filtered using information encoded in knowledge sources; in this inset, only findings related to "sign or symptom" are shown. (d) In this inset, only findings related to dexamethasone such as edema (to see whether the drug has an effect on reducing swelling) and cushingoid (a possible side effect of the drug) are shown. (e) Selecting a particular observation brings up details on demand such as related imaging studies and attributes extracted from the text.

3) *Details on demand*: Users can view patient data by interacting with the patient dashboard. Hovering over an observation brings up a tooltip with the original sentence(s) related to the finding. Selecting an observation opens a separate window [see Fig. 3(e)] that shows related patient data (e.g., slices of an imaging study that depicts the described finding) and other structured information (e.g., whether the finding is new or recurring, related blood test values).

As part of the iterative development process, we asked two physicians (a neuroradiologist and an oncology fellow) to view a set of patient cases and identify the strengths and weaknesses of the AdaptEHR interface. Initially, users were presented with an alphabetically sorted list of all findings identified in the patient record. However, both users found this list difficult to navigate given a large number of findings. As a result, we implemented functionality to filter this list by semantic type: in Fig. 3(c), only findings related to "signs or symptoms" are shown. Findings could also be filtered by relation: a user specifies target variables by selecting findings listed in the patient dashboard. Using properties such as the Markov blanket, related findings are identified based on the topology of the model and presented to the user, removing all other findings from view. Fig. 3(d) shows an example of filtering findings to those related to the drug dexamethasone: the original list of 80 findings is reduced to three by filtering irrelevant findings based on the model. Dates when the drug was prescribed are clearly delineated. Related findings, ranked by

strength of influence, are also shown to visually convey effects of the drug. Both users found the filters intuitive and effective in presenting relevant information, but situations arose when users disagreed with the information presented (or noted that an important piece of information was missing). This issue relates to the assumptions made when creating the disease model. Presently, while users can view the model, they cannot modify it. In addition, users can revert to a nonfiltered view at any time; a log of user interactions with the interface is being captured for future study. Finally, users found the summarization of trends in the patient record to be straightforward, yet both experienced some confusion when navigating a cluster of findings with multiple observations in a short period of time. Both users expressed that the details on demand view was particularly helpful.

VI. DISCUSSION

Given the diversity and amount of data being collected as part of the patient record, an efficient approach to retrieve and understand this information is crucial given the decrease in amount of time available for clinicians to perform tasks. We present a context-based framework for formalizing knowledge encoded in biomedical ontologies and disease models to support the integration and adaptation of information in the health record. While our study primarily utilizes ontologies and probabilistic disease models, another important source of domain knowledge is clinical practice guidelines. Computerized guidelines encoded in a guideline interchange format could be used to generate

personalized recommendations based on what is documented in the record.

Additional research is needed to characterize the information needs of physicians when performing different tasks. Cimino *et al.* [15] asked clinicians to review a set of patient records and found that they varied widely on the types of information identified as relevant. Development and incorporation of a cognitive model that considers user preferences when determining if a data element is relevant would be beneficial. As illustrated through our preliminary study, the current disease model may not fully meet the information needs of a physician performing a specific task. To address this issue, we will explore additional approaches such as incorporating published clinical guidelines and providing tools to modify the model. In addition, the health-care community must address the paucity of disease models that are available for incorporation into clinical decision support systems. While some models such as HEPAR II [30] have been disseminated, most domains do not have models that are public and vetted. Further development is needed to create robust tools for constructing, validating, and sharing models with the broader community. Formal evaluations are planned to further assess the capabilities of the structuring and mapping tools and the usability of the AdaptEHR interface. While we have obtained valuable feedback from presenting components of this system [28], [29], [31], a formal study is being conducted to compare the performance of physicians when using AdaptEHR versus existing means (e.g., manually generated slide presentation) while reviewing a case at a tumor board conference. Evaluation metrics will include time spent to prepare the case, ability to answer questions posed by board members (e.g., precision/recall of information retrieved from the patient record), and overall satisfaction with the interface.

ACKNOWLEDGMENT

The authors would like to acknowledge Dr. T. Cloughesy for providing access to cases seen at the University of California, Los Angeles (UCLA) Neuro-Oncology Clinic as well as the physician users and reviewers for their valuable feedback.

REFERENCES

- [1] R. Steinbrook, "Health care and the American Recovery and Reinvestment Act," *N. Engl. J. Med.*, vol. 360, no. 11, pp. 1057–1060, Mar. 12, 2009.
- [2] E. Abbo, Q. Zhang, M. Zelder, and E. Huang, "The increasing number of clinical items addressed during the time of adult primary care visits," *J. General Internal Med.*, vol. 23, no. 12, pp. 2058–2065, 2008.
- [3] C. Sicotte, J. L. Denis, P. Lehoux, and F. Champagne, "The computer-based patient record challenges towards timeless and spaceless medical practice," *J. Med. Syst.*, vol. 22, no. 4, pp. 237–256, 1998.
- [4] K. Davies, "The information-seeking behaviour of doctors: A review of the evidence," *Health Info. Libr. J.*, vol. 24, no. 2, pp. 78–94, 2007.
- [5] O. Ratib, "From multimodality digital imaging to multimedia patient record," *Comput. Med. Imag. Graph.*, vol. 18, no. 2, pp. 59–65, Mar./Apr. 1994.
- [6] R. E. Dayhoff, P. M. Kuzmak, G. Kirin, and S. Frank, "Providing a complete online multimedia patient record," in *Proc. AMIA Annu. Symp.*, 1999, pp. 241–245.
- [7] R. Santoni and E. Cellai, "Technical note: Multimedia clinical records: Results of a pilot project at the Radiation Therapy Department of Florence," *Br. J. Radiol.*, vol. 68, no. 808, pp. 413–420, Apr. 1995.
- [8] S. M. Powsner and E. R. Tufte, "Graphical summary of patient status," *Lancet*, vol. 344, no. 8919, pp. 386–389, Aug. 6, 1994.
- [9] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman, "LifeLines: Visualizing personal histories," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.: Common Ground*, Vancouver, BC, Canada, 1996, pp. 221–227.
- [10] 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 *Proc. 26th Annu. SIGCHI Conf. Human Factors Comput. Syst.*, Florence, Italy, 2008, pp. 457–466.
- [11] A. A. Bui, D. R. Aberle, and H. Kangarloo, "TimeLine: Visualizing integrated patient records," *IEEE Trans. Inf. Technol. Biomed.*, vol. 11, no. 4, pp. 462–473, Jul. 2007.
- [12] Q. Zeng and J. J. Cimino, "A knowledge-based, concept-oriented view generation system for clinical data," *J. Biomed. Inform.*, vol. 34, no. 2, pp. 112–128, Apr. 2001.
- [13] C. A. Morioka, S. El-Saden, W. Pope, J. Sayre, G. Duckwiler, F. Meng, A. Bui, and H. Kangarloo, "A methodology to integrate clinical data for the efficient assessment of brain-tumor patients," *Inform Health Soc. Care*, vol. 33, no. 1, pp. 55–68, Mar. 2008.
- [14] D. Sittig, A. Wright, S. Meltzer, L. Simonaitis, R. S. Evans, W. P. Nichol, J. Ash, and B. Middleton, "Comparison of clinical knowledge management capabilities of commercially-available and leading internally-developed electronic health records," *BMC Med. Informat. Decision Making*, vol. 11, no. 13, pp. 1–9, 2011.
- [15] J. J. Cimino, G. Elhanan, and Q. Zeng, "Supporting InfoButtons with terminological knowledge," in *Proc. AMIA Annu. Fall. Symp.*, 1997, pp. 528–532.
- [16] N. F. Noy, N. H. Shah, P. L. Whetzel, B. Dai, M. Dorf, N. Griffith, C. Jonquet, D. L. Rubin, M. A. Storey, and C. G. Chute, "BioPortal: Ontologies and integrated data resources at the click of a mouse," *Nucleic Acids Res.*, vol. 37, no. suppl. 2, pp. W170–W173, 2009.
- [17] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988.
- [18] P. J. Lucas, L. C. van der Gaag, and A. Abu-Hanna, "Bayesian networks in biomedicine and health-care," *Artif. Intell. Med.*, vol. 30, no. 3, pp. 201–214, Mar. 2004.
- [19] W. Hsu, A. A. Bui, R. K. Taira, and H. Kangarloo, "Integrating imaging and clinical data for decision support," in *Handbook of Research on Advanced Techniques in Diagnostic Imaging and Biomedical Applications*, T. Exarchos, P. A. and F. DI, Eds., Hershey, PA: IGI Global, 2009, ch. II, pp. 18–33.
- [20] E. W. Watt and A. A. T. Bui, "Evaluation of a dynamic Bayesian Belief network to predict osteoarthritic knee pain using data from the osteoarthritits initiative," in *Proc. AMIA. Annu. Symp.*, 2008, pp. 788–792.
- [21] J. R. Koiter, "Visualizing inference in Bayesian networks," Ph.D. dissertation, Dept. Comput. Sci., Delft Univ. Technol., Delft, the Netherlands, 2006.
- [22] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [23] M. Tong, J. Wu, S. Chae, A. Chern, W. Speier, W. Hsu, A. Bui, and R. Taira, "Computer-assisted systematic review and interactive visualization of published literature," presented at the Radiol. Soc. North America (RSNA) Annu. Meet., Chicago, IL, 2010.
- [24] GRADE Working Group, "Grading quality of evidence and strength of recommendations," *BMJ*, vol. 328, no. 7454, pp. 1490–1494, 2004.
- [25] A. A. T. Bui, J. D. N. Dionisio, C. A. Morioka, U. Sinha, R. K. Taira, and H. Kangarloo, "DataServer: An infrastructure to support evidence-based radiology," *Acad. Radiol.*, vol. 9, no. 6, pp. 670–678, 2002.
- [26] R. K. Taira, V. Bashyam, and H. Kangarloo, "A field theoretical approach to medical natural language processing," *IEEE Trans. Inf. Technol. Biomed.*, vol. 11, no. 4, pp. 364–375, Jul. 2007.
- [27] R. K. Taira, S. G. Soderland, and R. M. Jakobovits, "Automatic structuring of radiology free-text reports," *RadioGraphics*, vol. 21, no. 1, pp. 237–245, Jan./Feb. 2001.
- [28] W. Hsu, C. W. Arnold, and R. K. Taira, "A neuro-oncology workstation for structuring, modeling, and visualizing patient records," in *Proc. ACM Intl. Conf. Health Informatics (IHI)*, Washington, DC, 2010, pp. 837–840.
- [29] W. Hsu and R. K. Taira, "Tools for improving the characterization and visualization of changes in neuro-oncology patients," in *Proc. AMIA Annu. Symp.*, 2010, pp. 316–320.
- [30] A. Oniško, "Probabilistic causal models in medicine: Application to diagnosis of liver disorders," Ph.D. dissertation, Inst. Biocybern. Biomed. Eng., Polish Academy Sci., Warsaw, Poland, 2003.
- [31] W. Hsu, C. W. Arnold, A. Chern, R. K. Taira, and H. Kangarloo, "Tools for characterizing, visualizing, and comparing trends in neuro-oncology patients," presented at the Radiol. Soc. North Amer. (RSNA) Annu. Meet., Chicago, IL, 2010.