MEDINFO 2004
M. Fieschi et al. (Eds)
Amsterdam: IOS Press


# Automated Medical Problem List Generation: Towards a Patient TimeLine

**Alex A. T. Bui, Ricky K. Taira, Suzie El-Saden, Alicia Dordoni, Denise R. Aberle**

*Medical Informatics Group, University of California, Los Angeles, CA, USA*

**Abstract**

*The problem-oriented electronic medical record has been investigated as an alternative to source-oriented organization of patient data. At the core of a problem-oriented view is the medical problem list. Maintenance of the medical problem list is often manual, making it highly user dependent. We detail the beginnings of an automated medical problem list generator based on ICD-9: given a set of ICD-9 codes associated with a patient record, the system maps the codes (and related data) to an anatomy-centric hierarchy. 1 million patient encounters from an outpatient setting were reviewed to generate a unique set of 7,890 ICD-9 codes. Natural language processing of the ICD-9 string descriptions identified 1,981 anatomical terms, which were subsequently mapped to one of 21 anatomical categories. The output of the medical problem list generator was then used to create a problem-oriented, gestalt view of a patient's medical record. Preliminary evaluation of the generator revealed 100% recall, but only 60% precision. This initial work has highlighted several issues in defining a medical problem list, including questions of granularity and performance trade-offs.*

***Keywords:***

Problem-oriented medical records.

## Introduction

An essential element of evidence-based medicine is the use of external evidence in the form of recorded knowledge (*i.e.,* literature) from experts. Similar evidence can be obtained through consultation with expert colleagues at work. Currently, the majority of such communications involve direct physician-physician interaction: Safran reported that even in a hospital with a computer-based system, only 10% of "information transactions" occur using an electronic medical record (EMR), whereas 50% of such dealings involved face-to-face interactions between colleagues [1].

The ability to organize patient records according to medical problems (rather than traditional source-oriented methods) may better enable EMR-based information transactions between healthcare providers. With a patient's medical record re-organized around medical problems and visualizations of laboratory, imaging, and other procedures, a "gestalt" view of the patient's history can be created, explicitly detailing the supporting evidence for each problem (Figure 1). This kind of gestalt display can potentially facilitate communication and data sharing among physicians, not only in the same enterprise, but across multiple institutions.

Key to realizing this visualization paradigm is the ability to automatically generate a medical problem list, and from this list, instantiation of the display. This paper addresses some of the technical issues related to a medical problem list generator and its usage in creating this problem-centric EMR display.

## Materials and Methods

### Medical Problem List Generator

The input to the medical problem list generator is the contents of a given patient's electronic medical record (*e.g.,* reports, labs, imaging), with associated ICD-9 (International Classification of Diseases) codes. The output of the generator is an organized grouping of the data according to individual medical problems. The base requirements for the system are fourfold: 1) automatic creation of the problem list; 2) assignment of data to one or more problem categories; 3) no additional effort required by the physician (*e.g.,* no specific coding); and 4) favorable recall over precision, if compromises must be made. The latter concession is made, as if a specialist chooses only to view data linked to specific problem, it is important for all relevant documents to be included; some irrelevant or marginally related documents are presumed acceptable. Three major assumptions are made in this work:

1. Patient problems can be adequately summarized using ICD-9 code descriptions. Thus, the representation of a medical problem will be in the form of ICD-9 codes.
2. ICD-9 codes are available for each piece of medical data.
3. ICD-9 codes are accurately assigned to each medical encounter.

Our algorithm is dependent on accurate and consistent coding of medical data. The inputs and outputs of the generator are defined as follows:

- **Input:** Two vectors. One vector consists of string elements containing a list of new ICD-9 codes. The second vector contains a set of ICD-9 codes that represent the patient's current medical problem list.
- **Output:** One vector containing the new problem list. Each element in the vector represents a unique medical problem, and is itself a vector that holds references to the source data.

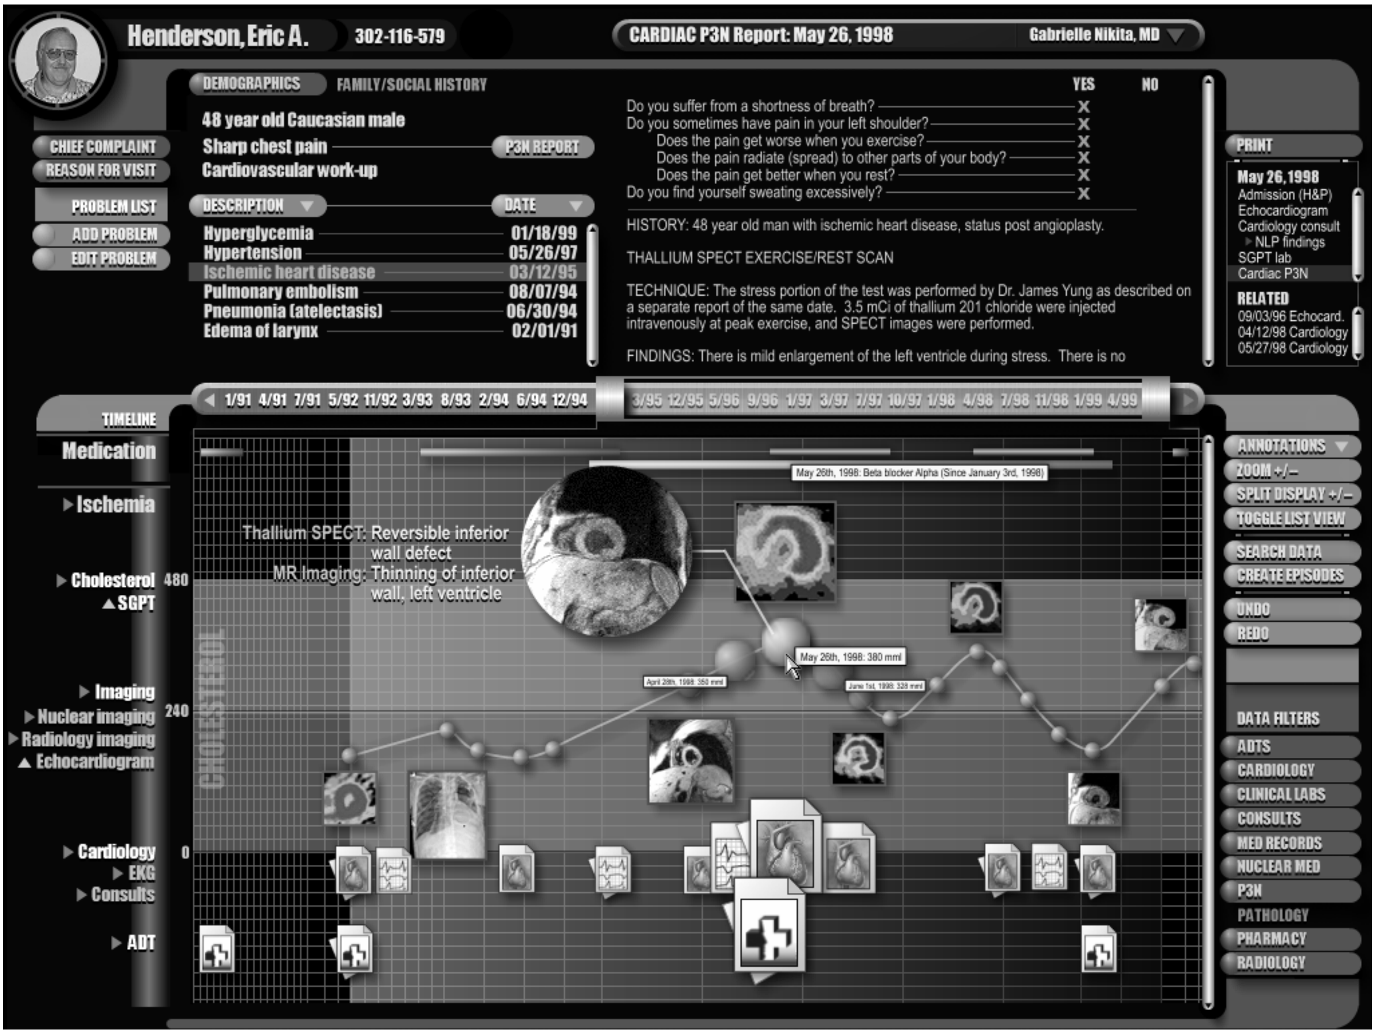


*Figure 1 - TimeLine visualization of medical problems in a patient with a history of chest pain. The initial cause of chest pain is shown in a chest x-ray demonstrating a pneumothorax, while myocardial perfusion is normal, as seen by a nuclear medicine study. Ischemic coronary heart disease, however, is the cause of later chest pain (26/5/98), as shown by an abnormal myocardial perfusion in the magnetic resonance imaging and nuclear medicine studies. Note the parallel tracking of cholesterol levels, with peak on 26/5/98*

Our first task was to exactly specify the allowed output classification states. The initial categorization of disease was intended to be indexed primarily via anatomy (*e.g.,* breast, brain, thorax, etc.), as medical specialization primarily follows this organization and because the majority of medical problems described by patients or physicians also have an anatomical relationship or reference. The anatomy-centric organization specified by the American College of Radiology (ACR) coding of disease was used [2]. Other categories, which could overlap with this anatomy-centric organization, were created to account for: systemic disease processes (*e.g.,* infectious diseases, endocrine disorders, hematopoietic disorders); psychiatric disorders; and drug and nutritional disorders. Our initial categories are summarized in Figure 2.

A lookup table was developed to find codes that have been automatically identified in one of our target 21 bins. Over 1 million patient encounter records were reviewed from an outpatient facility to compile a list of unique ICD-9 codes that were used; 7,890 unique codes were found. Natural language processing (NLP) methods were applied to each of these ICD-9 descriptions to automatically identify the topic anatomy [3, 4]. The NLP system extracted both anatomies of interest and anatomies to exclude from an ICD-9 description. For example, for the following ICD-9 code:

*ICD-9 code 569.81 - Fistula of intestine, excl rectum and anus*

the output of the NLP system is *includeAnatomy* = *intestine*; *excludeAnatomy* = *anus* and *rectum*. Using the list of 7,890 unique ICD-9 codes, a list of 1,981 unique anatomy descriptions was compiled. Each of these phrases was manually reviewed and assigned to one of the output classes (Figure 2), if applicable. Assignments were performed by a medical physician training in medical informatics (AD). This anatomy-to-disease category map further serves as a knowledge source for ICD-9 codifications outside of the initial list of 7,890 codes. Additionally, a list of unique disease terms was compiled from the list of 7,890 ICD-9 codes, and manually inspected to ascertain possible implied anatomy (e.g., from cardiomegaly, one can factor out the implied anatomy, heart).

| | |
|---|---|
| Breast | Infectious diseases |
| Brain and skull | Gastrointestinal disorders |
| Face, mastoids, neck (eyes, ear, nose, throat) | Genitourinary disorders |
| Spine and contents | Obstetrics/gynecology |
| Peripheral nerves | Endocrine |
| Mental disorders | Skin, hair, nails, subcutaneous tissues |
| Musculoskeletal (excluding skull, spine) | General signs and symptoms |
| Heart and great vessels | Chemical abuse and poisoning |
| Blood disorders | Laboratory |
| Vascular (excluding heart vessels) | Miscellaneous |
| Thoracic disorders | |

*Figure 2- Initial categories used to classify ICD-9 codes*

As a result of the above mapping – from ICD-9 description to anatomy to target problem list bin category – about 63% of the initial 7,890 ICD-9 codes could be automatically assigned to an output classification. The remaining unassigned codes were manually assigned to one or more of the target problem list bins. The final result of this process was a knowledge base (KB) allowing the original list of 7,890 ICD-9 codes to be simply indexed by this transformation table into one or more of the targeted problem list classes. For codes not included within the 7,890 manually verified entries of the KB, the following methods were used to assign target classes:

1. **ICD neoplasm cross-index.** Some codes related to neoplasms could be resolved by utilizing the ICD-9 CM neoplasm cross-index table [5]. For a given type of neoplasm, six different codes can be assigned, depending upon whether the neoplasm in question is malignant, benign, *in situ*, of uncertain behavior, or of unspecified nature. It was assumed that within a patient record, all specific neoplasms that are cross-referenced refer to the same phenomenon.
2. **Anatomy-index table.** ICD-9 codes not in the knowledge base are processed by our natural language processor, and anatomy and condition information are extracted. The anatomy description is searched for in the anatomy-to-problem list mapping described previously: if an entry is found, the corresponding problem category is assigned.
3. **Nearest match.** The original 7,890 ICD-9 codes that were mapped to problem categories are scattered throughout the space covered by ICD-9, having been compiled from over a million patient encounters. ICD-9 codes not found in this mapping are resolved by identifying the nearest cataloged ICD-9 code that does not cross sub-category boundaries. For instance, ICD-9 codes 001-009 cover the category, *Intestinal infection diseases*. The system will "interpolate" between these codes, such that an unknown code (*e.g.,* 004) is assumed to also be part of the *Intestinal infection diseases* category.

Figure 3 illustrates the basic algorithm employed on a patient record, using the developed ICD-9 knowledge base. The basic technique is as follows:

1. The set of ICD-9 codes associated with a patient's medical record is specified in a vector, sorted in reverse chronological order so that "newer" more specific findings are used as the basis for a medical problem grouping. Each ICD-9 code is then considered against a "working" problem list (WPL) that initially starts empty. Notably, the WPL may represent an existing medical problem list for the patient, and is also a vector.
2. For each new ICD-9 code, a comparison is made against the entries in the WPL, determining the "compatibility" of the new code *vs*. a code in the WPL. Compatibility is performed by comparing codes and using the lookup transformation table. If no equivalent code is found in the WPL, the new ICD-9 code is added as a new problem.
3. If the new code is found to be compatible with an existing WPL problem, a second question is posed: is the new problem more specific than the prior classification? For example, a general code of *cancer* may be further specified as *lung neoplasm*. If the code is not more specific, then only the initial date of problem onset is updated in the WPL. If the new ICD-9 code is more specific, then the WPL entry is replaced with the new ICD-9 code.

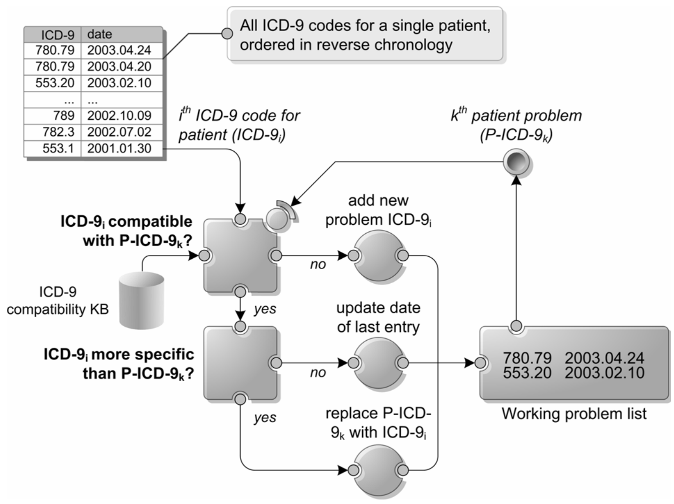


*Figure 3- Processing algorithm for creation of problem list. The system takes as input the set of ICD-9 codes associated with a patient record, and generates a "working" medical problem list*

**Medical Problem Visualization**

The second element of this research is focused on visualizing the medical problems, and is briefly described herein. Through use case modeling, we have developed a visualization paradigm, called *TimeLine* [6]. TimeLine provides a gestalt view for physicians and researchers of a patient's overall medical history, while still permitting for an expanded view of a specific medical problem with the supporting evidence and/or documentation pertinent to the problem (Figure 1). In point of fact, TimeLine is predicated upon an analytical perspective of medical care: much as the scientific method constitutes a methodical investigation of an unknown phenomenon, TimeLine attempts to detail the course of exploration and observations leading up to a definitive diagnosis and a treatment plan. Current data models are typically explanative, capturing what has been discovered with-out provision for the investigative aspect of the healthcare process and explicitly showing *how* a physician reached a conclusion.

Figure 4 shows an instance of the underlying data organization used by TimeLine: a *phenomenon-centric data model*. Firstly, each medical problem (*phenomenon*) may be complex, being associated with other lower level phenomena themselves. This recursive construct reflects a hierarchical view of medical problems, and the TimeLine view allows users to expand or collapse the view of a given problem and its sub-problems. Each problem has an associated set of (ab)normal *findings*, constituting anything "observed" about the phenomenon, including imaging, laboratory, reports, and other data generated from a healthcare encounter. Findings are used to relate effects to possible causes, and are evidence used to justify the final diagnosis. Associated with findings are: 1) *states*, which are snapshot descriptions of all relevant features of a finding at a single point in time; 2) *behaviors* reflecting changes between two states; and 3) *theories*, which attempt to provide a scientific explanation for the phenomenon/problem in question. To illustrate, Figure 4 shows the evidence at multiple finding levels that explain a patient's urinary incontinence: sagittal magnetic resonance (MR) images show meningomyelocele; a neurogenic bladder is shown in a voiding urethrocystogram; and a definitive diagnosis of bladder type IV (the reason for incontinence) is shown by video urodynamics.

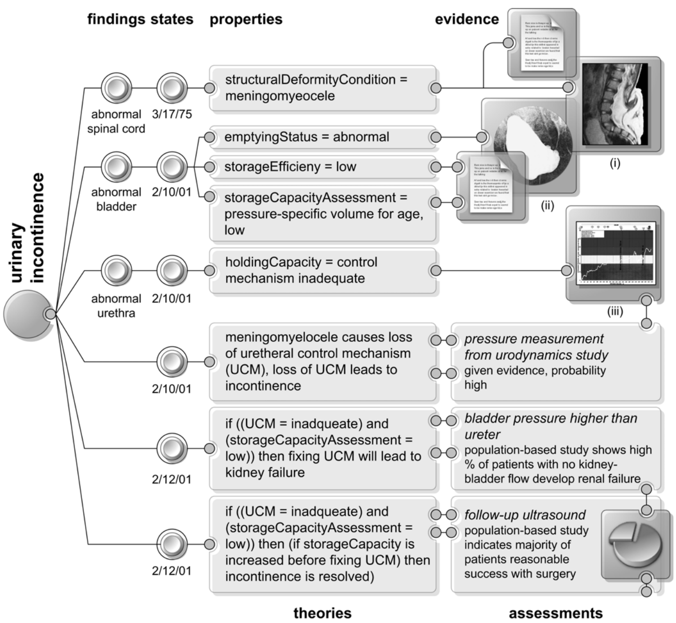


*Figure 4- Phenomenon-centric data model instance, demonstrating the relationship between evidence and patient presentation state for urinary incontinence. Meningomyelocele is demonstrated by a sagittal MR image (i), an abnormal bladder shape is documented by a voiding urethrocystogram (ii), and a video urodynamic study quantifies the patient's bladder capacity. The model is used to instantiate the TimeLine gestalt view.*

Problem list generation partly assists in instantiation of this phenomenon-centric model, helping TimeLine to classify different medical data elements into the problem-centric view.

## Results

We randomly selected 100 patients from a urology outpatient clinic. Data from the electronic medical record was retrieved, with all associated ICD-9 codes. The automatic medical problem list generator was given the set of ICD-9 codes per patient case, classifying the data into one of the 21 anatomical categories; large groupings of data were further sub-categorized by the system into sub-problems. The output of the generator was a hierarchical representation of the problem list, with each top-level node representing the most relevant medical problems (phenomenon) per patient record. Based on this output, the TimeLine generator used the retrieved medical data and the categorization to construct a web-based gestalt view for end-users. Average processing time per patient (independent of source data retrieval) for both problem list generation and creation of the web-based display took less than 20 seconds on a 2.8 GHz desktop personal computer.

A consensus group of three physicians and two medical informaticians reviewed the end categorization by the system, scoring each medical problem list entry for appropriate classification. This preliminary study showed that while recall was 100% (as would be expected, based on the design of the system), precision was only 60%; however, this conservative approach assured that all possible medical problems were captured.

## Discussion

The motivation for creating the problem list are summarized in [7], and include: 1) obtaining an overview of an unknown patient case; 2) updating one's knowledge about a known case; 3) searching for specific data related to a given hypothesis; and 4) solving clinical problems, including treatment planning. The specification and advantages of problem-oriented medical records (POMR) have been reported by many investigators over the years [8, 9, 10, 11, 12]. The benefits involve: improved communication between medical teams; improved precision/recall for information retrieval tasks; facilitating authoring of education material for students and patients; and data collection for clinical research (*e.g.,* patient cohorts).

The ability to visualize a medical problem list, in conjunction with the associated medical data (image, laboratory, etc.) detailing the evolution of a hypothesis to a final diagnosis, can be a major tool in improving communication among physicians, and promoting evidence-based research and education. In particular, lab findings and imaging are among the most objective types of evidence available for documenting a disease process and treatment response (*e.g.,* change in tumor size as a result of radiation therapy); rapid visualization of this information per medical problem can facilitate a quick understanding of a patient's overall health status. Indeed, [13] discusses how medical imaging is an effective method for communication not only between clinicians, but also with patients. As such, these two types of data play a predominant role in the TimeLine user interface.

The system has been implemented in a laboratory setting with satisfactory performance. It must, however, be evaluated further in the real-world environment to assess its practical utility. Major issues that need to be addressed beyond this preliminary work are:

- The complexities associated with the interaction of current and historical problems. Physicians are trained to distinguish when two (or more) incidents of a given medical problem are distinct (*e.g.,* a 5-year separation between skin rashes is probably two separate incidents); however, this knowledge is not explicitly modeled. When should the system split two occurrences of the same codified problem? Equally, the system also needs sufficient understanding to combine separately encoded issues that may in fact, during the course of investigation, turn out to be the same phenomenon (*e.g.,* a migraine headache and a brain tumor). This issue is intimately tied to the idea of episode creation, which can be highly dependent on user perspectives. Usage of additional information from the EMR (*e.g.,* CPT-4 (Clinical Procedure Terminology) codes, chief complaints, reason for visit) may be useful in providing further guidance to the system.
- The level of granularity needed to encode clinical details that describe the patient in a real-world setting. The initial choice of using an anatomy-centric organization may be too coarse: some users may find that the resultant classification of disease entities does not provide sufficient differentiation (*e.g.,* a urologist may find that having all urological conditions under a genitourinary grouping may be insufficient, and want further detail). Again, this matter is probably driven by user preferences, as different types of physicians (*e.g.,* specialists *vs.* primary care physician) will require varying levels of detail.
- The performance balance between recall and precision. What are the consequences of decreasing recall performance to improve precision? From a purely technical level, doing so would decrease the amount of "non-relevant" data classified per problem, at the risk of excluding important information – this many not be acceptable if physicians solely rely upon the gestalt view to make clinical decisions. Conversely, what is the ultimate impact of having high recall and lower precision? Physicians will then be faced with possibly navigating through too much information, decreasing the overall utility of the classification mechanism.

### Acknowledgments

This work was supported in part by grants from the National Institute of Biomedical Imaging and Bioengineering (NIBIB) PO1-EB00216, RO1-EB000362, and from the National Library of Medicine (NLM), T15-LM07356. The authors would also like to thank Dr. Hooshang Kangarloo for his insight and help in this research.

## References

[1] Safran C, Sands DZ, Rind DM. On-line medical records: A decade of experience. *Methods Info Med* 1999, 38:308-12.

[2] American College of Radiology. *Index for Radiological Diagnoses.* 4th Edition. American College of Radiology. 1891 Preston White Dr. Reston, VA 22091, 1992

[3] Taira RK, Soderland S. A statistical natural language processor for medical reports. *Proc AMIA*, 1999; pp. 970-974.

[4] Taira RK. Semantic phrasal chunking of medical expressions in free text reports. In *MedInfo 2004*.

[5] *International Classification of Diseases. 9th Revision Clinical Modification*, 5th edition, Hospital Edition, Volumes 1, 2, 3 PMIC, 1998; pp1018-1048.

[6] Bui AAT, Taira RK, Churchill B, Kangarloo H. Integrated visualization of problem-centric urological patient records. In Valafar F, ed., Techniques in Bioinformatics and Medical Informatics, *Annals of the New York Academy of Sciences*, 2002; 980:267-277.

[7] Nygren E, Henriksson P. Reading the medical record, Part I: Analysis of physicians' ways of reading the medical record. *Comp. Methods and Programs in Biomed*, 1992; 39:1-12.

[8] Weed LL. Medical records that guide and teach. *N Engl J Med.*, 1968; 278:593-600.

[9] Bolens M, Borst F, Scherrer JR. Organizing the clinical data in the medical record. *MD Computing*, 1992; 9:149-155.

[10]Poon AD, Fagan LM, Shortliffe E. The PEN-Ivory Project: Exploring user-interface design for the selection of items from large controlled vocabularies of medicine. *J. Am Med Inform Assoc*, 1996; 3:168-183.

[11]Zeng Q, Cimino JJ, Zou KH. Providing concept-oriented views for clinical data using a knowledge-based system: An evaluation. *J. Am Med Inform Assoc*, 2002; 9(3):294-305.

[12]Institute of Medicine. Key capabilities of an electronic health record system. Letter Report. National Academy Press, Washington DC 2003.

[13]Kaplan B, Shaw NT. People, organization, and social issues: Evaluation as an exemplar. In *Yearbook of Medical Informatics*, 2002; pp. 91-99.

### Address for correspondence

Alex Bui, PhD

UCLA Medical Informatics Group & Biomedical Informatics Center

924 Westwood Blvd Suite 420

Los Angeles, California USA 90024

E-mail: buia@itmedicine.net