Concept-Based Electronic Health Records: Opportunities and Challenges

Shahram Ebadollahi[*§], Shih-Fu Chang[†], Tanveer Syeda-Mahmood[‡], Anni R. Coden[*], Arnon Amir[‡], Michael A. Tanenblatt[*]

[*] IBM T.J. Watson Research Center, Hawthorne, NY

[†] Digital Video Multimedia Lab, Columbia University, NY

[‡] IBM Almaden Research Center, San Jose, CA

[§] E-mail: ebad@us.ibm.com

ABSTRACT

Healthcare is a data-rich but information-poor domain. Terabytes of multimedia medical data are being generated on a monthly basis in a typical healthcare organization in order to document patients' health status and care process. Government and health-related organizations are pushing for fully electronic, cross-institution, integrated Electronic Health Records to provide a better, cost effective and more complete access to this data. However, provision of efficient access to the content of such records for timely and decision-enabling information extraction will not be available. Such a capability is essential for providing efficient decision support and objective evidence to clinicians. In addition researchers, medical students, patients, and payers could also benefit from it. We present the idea of concept-based multimedia health records, which aims at organizing the health records at the information level. We will explore the opportunities and possibilities that such an organization will provide, what role the field of multimedia content management could play to materialize this type of health record organization, and what the challenges will be in the quest for realizing the idea.

We believe that the field of multimedia can play a very active role in taking healthcare information systems to the next level by facilitating the access to decision-enabling information for different types of users in healthcare. Our goal is to share with the community our thoughts on where the field of multimedia content management research should be focusing its attention to have a fundamental impact on the practice of medicine.

Categories and Subject Descriptors

H.3.1 Content Analysis and Indexing, H.4.2 Types of Systems (Decision Support), H.5.1 Multimedia Information Systems

General Terms

Algorithms, Management, Design

Keywords

Medical Decision Support, Multimedia Analytics, Concept Detection, Electronic Health Records

1. INTRODUCTION

1.1 National Initiatives – Need for Electronic Health Records

The business of healthcare is experiencing a major overhaul regarding the role of information technology in everyday clinical practice. For the first time, the department of health has created the "Office of National Coordinator of Health Information Technology" (ONCHIT). There is a new mandate on the part of the government and health-related organizations to create and adopt "Electronic Health Records" (EHR), which allows providers and healthcare organizations to electronically store and exchange health-related information to anywhere a patient needs care.

Multiple reasons and requirements necessitate this move on the part of the government as stated by the head of ONCHIT [1]:

- "Informed clinical practice" - Despite tremendous amounts of data being generated and captured on a daily basis in the healthcare organizations, clinicians and other practitioners are deprived of decision-enabling information throughout the process of care. This problem negatively impacts the workflow in the healthcare center and more importantly could result in medical errors [2].
- "Interconnection of clinics" – People seeking healthcare due to many different factors visit several different institutions in their lifetime and also change their primary care physicians. Because of that in the current model for healthcare there is a lot of duplication in data acquisitions and also lost medical records. The goal of EHR is to overcome the single institution boundaries and allow the medical records to follow the patient wherever he/she seeks healthcare services. Ideally there should be single universally-accessible patient medical record, although different pieces of it could be physically residing in different locations.
- "Improved population studies" – When the multi-institution electronic health record becomes available, it could potentially be used to monitor public health,

- disease outbreaks, issues regarding national security and control the quality of care.

1.2 The Role of Multimedia in Healthcare

Healthcare information technology, which has traditionally been a text-based domain, has become ever increasingly dependent on a large array of multimedia data. This in major part is due to the advancements made in technologies related to diagnostic imaging devices, which have resulted both in new kinds of modalities of imaging and cheaper acquisition devices. In a typical hospital terabytes of data are being generated each month. The multimedia data in a hospital could be classified as following:

- Text: lab reports, clinical notes, pathology reports, etc.
- Image/Video: various modalities of medical imaging each providing a different and complementary view into the structure and function of patient's organs. Examples are: magnetic resonance, functional magnetic resonance, ultrasound, computed tomography, positron emission tomography, microscopy, pathology, etc. [3, 4]
- Other: signals such as electro-cardiogram, electro-encephalogram, etc.

Several different information systems [5, 6] are being employed today in healthcare centers for managing the array of medical data as mentioned above. Hospital Information Systems (HIS) are used for managing lab reports, electro-cardiograms, medical consults, etc.; Radiological Information Systems (RIS) are employed to manage diagnosis reports generated from reviewing medical diagnostic images, and Picture Archiving and Communication System (PACS) manages medical diagnostic images. The protocols for storing and communicating such data are specified by standards such as HL7 [7] and DICOM [8]. The US government has launched a large scale effort to create Regional Health Information Organizations (RHIO) [9] that will share medical information across hospitals and health-care institutions allowing mining of information while addressing privacy and rights issues.

1.3 Beyond Electronic Health Records – An Information-Centric Paradigm

To date, medical records are fragmented into multiple different systems. There are systems that aim at aggregating all patient-related data across RIS/HIS/PACS, which are being referred to as the Electronic Medical Records (EMR). EHR, mentioned earlier, is the extension of EMR to multiple institutions, and aims at not only collecting patient records across disparate systems, but also across geography and institutions. However, users of the EHR need to interact with the data at a file or document level. That is, although EHR helps with data aggregation it does not necessarily provide the means for easier access to decision-enabling information, which are embedded in the content of each piece of data it holds. Current medical record organization is not designed to be readily amenable to providing efficient access to the elements of the content of the data.

Every diagnostic imaging modality, every type of experiment and lab result, and every textual piece of data residing in the EHR, documents the attributes, characteristics and behaviors of a number of concepts. Concepts are anatomical and sub-anatomical units of the body, diseases, medical conditions, etc. For example, an echocardiogram video is used to provide visual information about the structure and function of the heart from different view angles [10]. The video document itself is not what interests a clinician or other users, rather the attribute and behavior of the sub-organs of the heart, such as valves, ventricles, and myocardium that indicate the medical condition of the heart are what the physician wants to be informed about. The text report of the echocardiogram summarizes the condition of those sub-organs and indicates the most probable diagnosis. Organizing the patient medical record in a way that the clinician or any other user could access all information extracted from different sources regarding any single concept of interest (e.g. "mitral valve" of the heart) could result in advantages that are not possible to obtain with the existing paradigm.

Now that the government and the healthcare institutions are pursuing EHR and the fully electronic patient records, we believe a new paradigm is required to help the domain of healthcare to better and more efficiently utilize the vast amount of data captured in the medical records. In this paper, we introduce the "concept-based multimedia medical records", which aims at aggregating snippets of information extracted from multiple sources of multimedia data for all concepts in the context of patient's health status.

We believe that the field of multimedia can play a very active role in a joint effort to take the "information poor" field of healthcare to the next level by facilitating the access to decision-enabling information to physicians, patients, researchers and payers. Our goal in this paper is to share with the community where we believe the field of multimedia content management should be focusing its attention in order to have a fundamental impact on the practice of medicine. We suggest the concept-based organization paradigm as a natural avenue for integrating multimedia with health care information of other modalities, thereby facilitating the maximum benefits of using multimedia in health care.

In the rest of the paper we first elaborate on the idea of concept-based health records. Then we will discuss the requirements for materializing it, followed by the opportunities and capabilities it brings about along with a number of challenges that need to be met.

2. CONCEPT-BASED HEALTH RECORDS ORGANIZATION

Throughout a person's clinical history, different types of media data are being captured and stored in healthcare centers. Together they provide a "medical representation" of that person. The purpose of each piece of data in such records, regardless of its modality, is to document certain information about some aspect of the patient's health status. For example, an echocardiogram video is acquired to provide information about the structure and function of the heart and to reveal possible problems with the sub-organs of the heart, such as the left ventricle or mitral valve. In parallel, an electrocardiogram, which documents the electrical signals stimulating the heart, is also captured and used by the clinician to extract certain information about the problems of the sub-organs of the heart, such as atrial fibrillation which involves the atria and is caused by irregularities in the impulses generated by the sinoatrial node. The array of captured data provides complementary information about the patient, like the above example that one can derive some information about the atria of the heart from the echocardiogram and some from the electrocardiogram.

Current systems for managing medical records deal with the patients' data at the document level. For example, diagnostic images such as echocardiograms are being stored in a PACS [5, 6] system according to the DICOM [8] standard, which attributes certain meta-data to the document, such as patient ID, ordering physician, date, etc. If a clinician wants for example to know the prior condition of the right atrium of the patient's heart, she first retrieves the relevant diagnosis report of the patient and look up what was reported on the status of the right atrium. If she wants to see the function of the right atrium first-hand, she needs to retrieve the corresponding echocardiogram image sequence, browse it to find an appropriate "view" of the heart and then see how the right atrium behaves. The situation becomes even more time consuming when the clinician needs to study and compare side-by-side the function of the "right atrium" of the heart between two different echocardiograms taken at different times in order to assess the efficacy of the treatment plan for example.

2.1 INFORMATION NEEDS

There are different categories of users of the medical records at different levels of expertise and with different information needs and tasks. These are physicians, researchers, medical students, payers, and patients. In this sub-section our goal is to provide some example scenarios of the information needs in each category and how the medical record is used today to cater to those needs, what its deficiencies are and what needs to be done differently for better information provision to different types of users.

2.1.1 Physicians

Bui et al [11] describe the ideal process of diagnosis as one consisting of the four stages of "patient assessment"; "conducting diagnostic procedures", such as ordering lab tests or acquisition of diagnostic images; "initial analysis"; and "final diagnosis". At each point during this process physician could have different information needs to be able to better conduct the diagnosis process.

Conducting differential study is common in assessing the efficacy of the treatment plan or understanding the progress of a disease. For this, the clinician wants to see the change in an organ or a sub-organ and its properties and characteristics over time. For example, a cardiologist might want to see the status of the boundaries of the left ventricle of the patient's heart as captured by two echocardiograms taken at different times to assess the extension of a myocardial infarction problem. Currently, the cardiologist needs to retrieve the right views of the heart from the two echocardiogram stored in the PACS system, and view them side-by-side to do his assessment. Ideally, she should be able to query the medical records for left ventricle and see all information pertaining to her needs and task aggregated for the concept of left ventricle.

Cardiologists routinely need to measure certain attributes of the cardiac sub-organs, such as the "ejection fraction" of the "left ventricle" of the heart, which indicates the pumping capacity of the heart. Currently, while reviewing the echocardiograms such measurements are being made. In the ideal situation, the system should be able to automatically know how to locate the "left ventricle" and then apply certain quantification algorithms to it in order to obtain such measurements and indicators to be provided to the cardiologist at the time she reviews the case.

Decision support and obtaining evidence is another service that clinicians could benefit from. At any point during the diagnosis process, a clinician might have certain information needs. Smith [12] suggests that the most frequently consulted sources of information for clinicians in this context are other specialists, text books and journal articles, in that order, followed by others. The most frequently consulted information source by clinicians, i.e. specialists, possess a vast amount of knowledge accumulated throughout years of practice, which is hard if not impossible to describe and capture by an external knowledge base such as books or articles. However, such knowledge is reflected in the patients' medical records in the form of diagnosis reports and interpretations of patients' conditions. Therefore, by using the combination of the medical records, and external sources one can potentially provide answers to the clinicians information needs. For a specialist such information needs usually occur in the context of rare medical cases or the ones that are difficult to be readily diagnosed. Smith [12] also mentions that clinicians occasionally also look for affirmation and feedback on the actions they plan to take during the process of diagnosis.

Linking the information snippets embedded in the multimodal medical records, such as the visual manifestation of myocardial infarction in the echocardiogram, to a part of the diagnosis report providing some information about it and to a passage in a textbook or the abstract of an article reporting the latest treatment for this ailment is almost impossible with the current medical record organization system. In another scenario, a novice clinician might need some evidence to support his decision about diagnosing a case with myocardial infarction. In this scenario the system should be able to know where the left ventricle of the heart is in the echocardiogram and submit a query based on the motion pattern of the activity of the walls of the left ventricle, retrieve all cases with similar pattern of motion and provide to the clinician the summary of their associated diagnosis reports. Such a system requires being able to access the visual manifestation of the concept of left ventricle in the echocardiogram.

2.1.2 Students

Students under medical training frequently consult teaching files, which have examples of medical cases. The reason for consulting such collections of medical records is to educate them about the different information related to a condition, how the condition manifests itself in a diagnostic image, what are its attributes in the lab results, or how it was diagnoses and what treatment plans were considered for it. They also consult articles in medical journals and textbooks to find out more about that information. It would be very beneficial for the medical students to have access to a multimedia summary of the information of interest to them that could allow them to navigate through the information. For example, while they are studying about mitral valve regurgitation, they can see the selected snippet of the diagnosis report pointing to that problem, while at the same time seeing an illustrative view of the heart from the echocardiogram that clearly shows that abnormality, where mitral valve has been labeled. This information could also be linked to a passage in a textbook that explains mitral valve regurgitation and how it could be detected using echocardiogram. Such kinds of automatic linkage between the information snippets embedded in the different pieces of data is lacking from the current medical record organization.

2.1.3 Patients and Payers

These days, more patients want to be informed about their own medical conditions. People usually consult online information sources to educate themselves about etiology, treatment options,

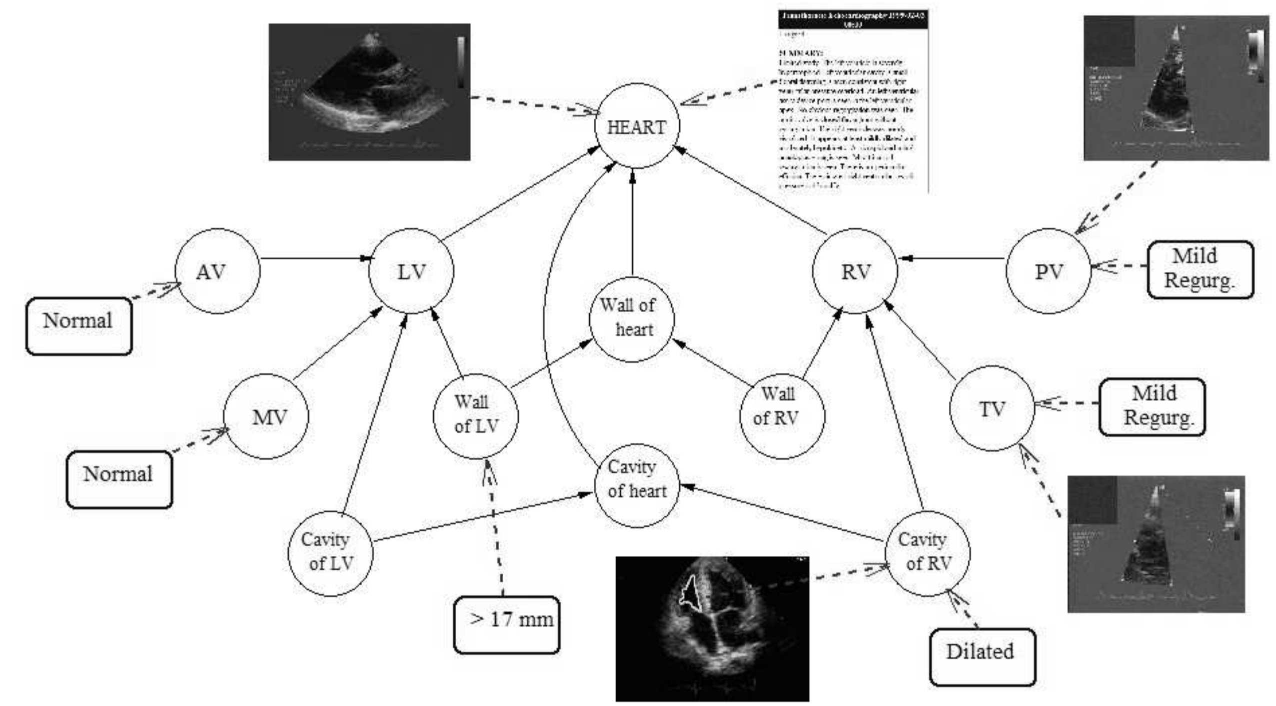

Figure 1. Partial ontology with concepts linked to their multimedia manifestations. (RV=Right Ventricle, LV=Left Ventricle, AV= Aortic Valve, MV = Mitral Valve, PV = Pulmonary Valve, TV = Tricuspid Valve)

and the prognosis of the medical conditions they have been diagnosed with. They could benefit from getting access to their own records in an illustrated and explained version, so they can better understand their medical conditions.

Payers, such as insurance companies, on the other hand need to monitor the use of resources to minimize costs, control risks, and provide the best service level. This requires verification of diagnosis procedures, referrals and tracking of treatment results. Using concept-based analysis of the patient's EHR it would be possible to compare the content of lab tests with their results, the diagnosis and the suggested treatment, prevent misuse of expensive tests (e.g., MRI) when not justified by results of earlier, less expensive or less obtrusive tests (e.g., EKG).

2.2 A NEW PARADIGM: CONCEPT-BASED ORGANIZATION OF MEDICAL RECORDS

There is a mismatch between the information needs of the clinicians and other users of the medical records and the way the medical records are being organized. It results in a decreased utilization of the medical records, inefficiency, and lack of decision-enabling information for users. The medical record organization paradigm should evolve from a document-level record management to an information-level approach. We propose the "concept-based health record", which aims at providing access to the manifestations of various concepts in the context of care process. Figure 2 shows the framework for creating such record organization. As shown in the figure, documents residing in the EHR are processed by appropriate analytical engines (AE) guided by relevant domain knowledge. For example an AE could be one that specializes in annotating and locating cardiac chambers in a given view of the echocardiogram video, or one that can parse the text of the diagnosis report and identify the parts of the text explaining or providing information about a given concept.

The function of the AEs and the concepts they aim at locating their manifestations in the data, highly depends on the definition of the concepts. Also after concept identification, one should know the relationships between them in order to organize them in a coherent fashion. Hence, concept-based organization of the medical records requires the definition of concepts and their relationships. In other words, the structure of knowledge pertaining to the clinical practice and patients' healthcare should be known. There are ontologies available in the domain of biomedicine, which capture and represent the concepts and their relationships in that domain. The most notable, multi-purpose, and widely uses is the National Library of Medicine's "Unified Medical Language System" (UMLS) [13], which integrates a large number of distinct source terminologies. The concepts are categorized in several semantic categories such as: organisms, anatomical structures, biological functions, chemicals, events, physical objects, etc. The links between different concepts indicate important relationships, which are of the following nature: "is a", "physically related to", "spatially related to", "temporally related to", "functionally related to", and "conceptually related to". C. Rosse, J.F. Brinkley and their colleagues in the University of Washington [14, 15] have argued that UMLS lacks the required granularity, adequate semantic types and relationships to comprehensively represent anatomical concepts and hence created the foundational model of anatomy as an enhancement to UMLS. They also have argued that all medical problems are grounded in anatomical concept. Medical diagnostic

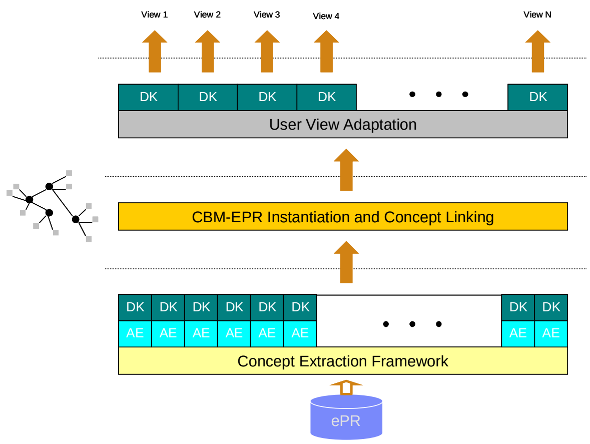

Figure 2. Architecture for concept-based medical record organization

images provide information about the anatomical units of the body and hence the appropriateness of anatomical concepts for representing clinical information [15].

Figure 1 illustrates an example of a part of a "concept-based record" of a patient pertaining to her echocardiogram examination. This example is based on a sub-section of the "Foundational Model of Anatomy" in the context of the sub-organ right atrium. This kind of medical record organization can also help in visualizing information across anatomical resolution axis. For example, one may traverse the records from the coarse organ level all the way to cellular level, where the concepts at each level are properly populated by the information extracted from documents pertaining to the right resolution. Information snippets extracted from molecular imaging, which aims at imaging molecular pathways in vivo, could be extracted and associated with the concepts at the molecular level, while at lower resolution levels one would associate information snippets extracted from echocardiograms or magnetic resonance images or other non-image sources to the relevant concepts. The linkage between the concepts along the resolution axis are captured in the ontology, which allows the user or the automatic processing to easily traverse such structures in order to aggregate the required information.

On the other hand, at any given time and any new clinical encounter, new pieces of information could be added to the patient concept-based health record. Such a health record resembles a growing attributed graph structure, where the nodes are the concept and the links between them illustrate concept relationships and to each node certain attributes are associated, which embody the different manifestations of that concept in different sources of data. **Error! Reference source not found.** illustrates the evolution of the patient's concept-based medical record with time. On the left-hand side two documents are generated ("findings report" and an "echocardiogram") and their elements of content are identified and organized according to the available semantic knowledge. At a future instance of time, a new piece of document is captured. Certain information snippets extracted from its content just provide different visual manifestations of the concept, whereas some new concepts are also added to the concept-based records to reflect the additional concept manifestations.

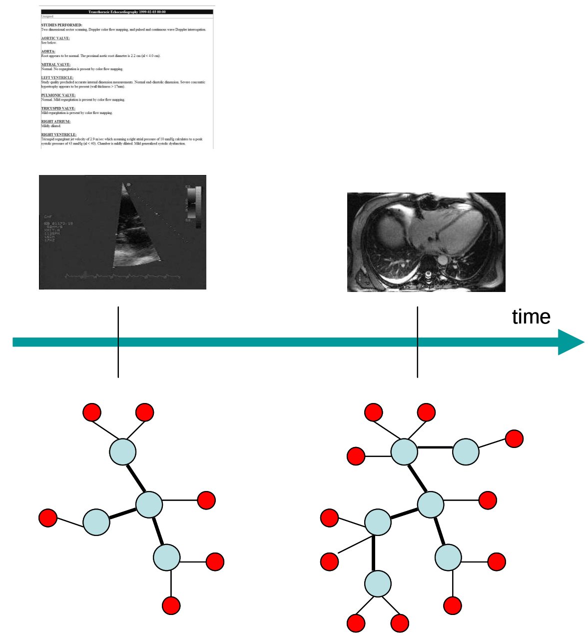

Figure 3. Addition of new concepts and their related information snippets to the concept-based health record

Due to different concepts of interest in different domains of clinical practice, one can also imagine having a federation of ontologies, where each is used appropriately according to the clinical procedures that the patient experiences. The individual ontologies could also be linked via certain common concepts.

We should note here that the concept-based electronic health record could provide a unifying application framework to all the multimedia research activity that has been done and is currently under way in the various component fields. It is not proposing a new way of analyzing the data, but a novel way of organizing the health records, which is more intuitive and hopefully more useful for supporting various information needs with respect to the medical records.

3. MULTIMEDIA CONTENT STRUCTURING FOR ENABLING CONCEPT-BASED ELECTRONIC HEALTH RECORDS

The main and foremost technology required for the materialization of the concept-based health records is multimedia "content structuring". The majority of the medical data residing in the electronic health records are unstructured or in other words in a raw format. Structure, generally refers to distinct semantic units of the content of the data. Structuring a multimedia document requires understanding the content and being able to associate tags with different elements of the content. In the context of the concept-based electronic health records, content structuring is equivalent to identifying the multimedia manifestations of the concepts pertaining to a given patient's medical records.

Content structuring has been the focus of research in different domains for a long time. For diagnostic images, structuring is equivalent to employing and devising appropriate computer vision techniques that can segment, localize and label the content of the images. Medical diagnostic image segmentation has a long history in computer vision and still is being considered an important area of research. Various approaches, some generic and some specialized to a particular modality of imaging have been developed throughout the years. For surveys of the model-based and atlas-based methods for medical image segmentation see [16, 17]. For video data in addition to the structuring activity required for images, one should also breakdown the video into its constituent temporal units of structure [18]. In case study (1), coming later in this section, we will provide a brief overview of our prior work on automatic structuring of the spatio-temporal content of echocardiogram videos. For textual data, such structuring requires applying the techniques of natural language processing in order to identify the syntax and semantics of the elements of the text. Chen et al. [19] provide a good overview of related NLP techniques for textual content structuring. Case study (2) below, provides our recent work on structuring pathology reports in the context of cancer care. For other types of multimedia data employed in the process of care, content structuring approaches have also been devised. Fernandez and colleagues [20] report their approach for identifying and tagging “R-peaks” in electrocardiograms.

In the following we briefly present two case studies which briefly explain some of our related work in addressing content structuring for echocardiogram videos and for textual Pathology reports in the domain of cancer care.

3.1 CASE STUDY: ECHOCARIOGRAM VIDEO SPATIO-TEMPORAL CONTENT STRUCTURING

Echocardiography [10] is a diagnostic imaging modality that uses ultrasound to capture the structure and function of the heart. A comprehensive evaluation entails imaging the heart in several planes (aspects) by placing the ultrasound transducer at various locations on the patient's chest wall. The recorded image sequence, echocardiogram video, displays the three-dimensional heart from a sequence of different two-dimensional cross sections (views). Due to the low cost and non-invasive nature, echocardiograms are a popular mode of imaging the heart in the domain of cardiology.

From the video analysis point of view, echocardiograms have inherent spatio-temporal structure. Recovering their elements of structure, which are the visual manifestations of the cardiology-related concepts, is important in order to better characterize their content by targeting feature extraction tools to targeted areas of interest and also for linking the visual elements of the content to external information such as text snippets extracted from the diagnosis reports.

Figure 4 below illustrates the system used for automatic content structuring of the echocardiograms. The system decomposes the echo video into its temporal units of structure, which are views, modes, and cardiac cycles. It also tags the different cardiac chambers in the different views of the echocardiogram video.

For view identification and labeling the cardiac chambers of the echocardiogram video, in [21] we employed undirected graphical models to learn the spatial arrangement of the cardiac chambers in the different views of the echocardiogram. Those models were then used for automatic view recognition in a discriminative framework. As a by product of view recognition we also labeled each cardiac chamber. Using the electrocardiogram signal embedded in the echo video, and also using the shape of the region of interest we decomposed the echo video into its modes and cycles [18].

The overall system was used in the context of PERSIVAL [22], which aimed at providing a personalized multimedia summary of multiple information sources for different types of users.

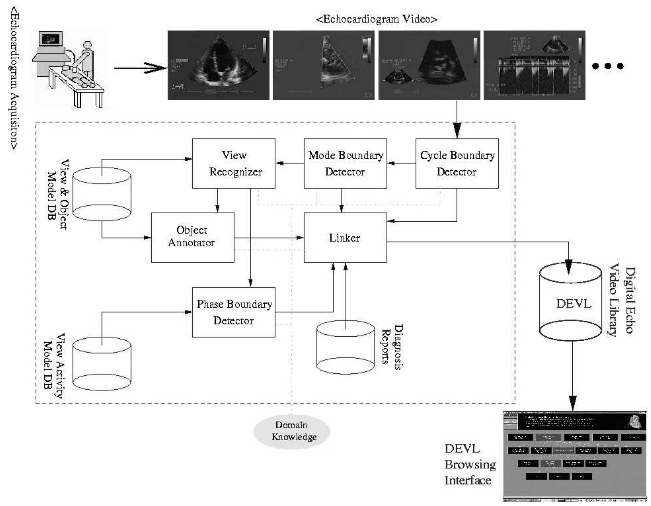

Figure 4. System for spatio-temporal structuring of echocardiogram videos

3.2 CASE STUDY: TEXT ANALYTICS FOR CANCER PATHOLOGY REPORTS

Recent advances in medicine showed the link between gene mutations and certain forms of cancer. For instance, researchers found a link between gene mutations and incidences of breast cancer. They have also shown that prophylactic surgery decreased the incidence of breast and ovarian cancers [23]. These latest results are based on a study following several hundreds of patients over 40 months. The effort to find an appropriate cohort of patients for such a study is quite time consuming, and one of the applications of MedTAS (Medical Text Analysis System) is to facilitate this process. Information about patients is collected in different places within a medical institution, and the majority of the information is in structured format, likely in multiple databases. Simple examples of structured information are demographic information and billing information. However, detailed patient information is recorded in unstructured text format, both in clinical notes, written by the attending physicians, and pathology reports, compiled after a biopsy or surgery Pathology reports are the definitive source for a cancer diagnosis, and hence mapping the information from the unstructured format of a pathology report into a structured and interpreted format is a critical tool.

Towards this end, a *case model* (Figure 5) and tools to populate the model were developed for MedTAS. This model tries to conceptually separate the different pieces of knowledge needed to describe the disease, evaluate procedures and correlate it to other parts of information stored within the institution. The *case model* defines four sub-models: the *tumor model*, the *specimen model*, the *lymph node model* and the *disease model*. The core model is the *tumor model*, which can either refer to a primary or a metastatic tumor. A *tumor model* has many characteristics, such as diagnosis, site, grade, invasion and size, but clearly can be augmented to accommodate others. A *specimen model* contains one or more *tumor models*, as well as information about the procedure used to obtain the specimen, such as surgery or biopsy, plus surgical margins. A complete disease diagnosis is based on all the *tumor models* and, in addition, on the information about the lymph nodes, which are objects with characteristics such as the number of positive nodes. Furthermore, the *disease model* contains other information, such as gross description and the staging information, critical in determining the treatment plan. The *case model* consists of the disease model and links to the tissue bank.

MedTAS (Medical Text Analysis System) automatically populates the cancer disease model. MedTAS is build on top of IBM's open source Unstructured Information Management Architecture (UIMA) which is currently available from SourceForge. MedTAS is a set of modular components (called annotators) that are applied in sequence to a collection of input documents. Some of these annotators are institution specific, whereas others are more general and can be applied in many different settings. We will briefly describe the sequence at a high level. The first and quite crucial realization is that the document structure contains a wealth of implicit information which has to be exploited to correctly extract the desired information and resolve co-reference. For example, the characteristics of a single tumor are in general described in several sections of a pathology report and although it is clear to a human which sections describe the *same* tumor, it is not trivial to teach a program the "obvious". MedTAS maps the pathology report into a virtual document format, where labels/titles/related sections are resolved. All subsequent processing is based on this virtual document format. Cleary this mapping is institution dependent, as there is no standard for writing pathology reports.

The next step, and the key to populating the cancer disease model, is named entity recognition. Since one of the ultimate goals of this process is to provide for subsequent information retrieval, mapping the named entities into a controlled vocabulary is critical. Fortunately, we can use ICD-O-3 (International Classification Disease – Oncology, version 3), which is a standard, published terminology that describes both topography and diagnosis. Other classifications that provide similar coverage could easily be substituted for ICD-O. Additionally, we developed a set of rules to create additional synonyms for the standard ICD-O entries. The advantage of using rules over manually creating synonyms is that they guarantee the correctness of the synonyms, as well as complete coverage. Our terminology lookup methodology allows for the addition of custom terminologies (and it is very simple to add them). Terminology lookup is word-order independent and discontinuous lookup based on semantic classes is enabled. We have done some preliminary work on non-repeating head nouns and list processing on which we will report in the future.

One other crucial aspect in this domain is negation. Our negation detector is based on natural language understanding, in particular in determining the correct noun phrases, which implies detecting correct sentences boundaries and part of speech tags. This is a challenging task in pathology reports, since they are mostly written in sentence fragments. We are in the process of adding higher level concepts (object, subject etc) to correctly identify the scope of the negation in conjunctive sentences.

Certain cancer characteristics require specialized processing for characteristics such as grade. Grading schemata differ by body parts, and each may be expressed differently and with different parameters. Furthermore, a pathologist interpreting a tissue slide cannot always clearly map what they see into the codified grading schema. We are exploring linking the tissue images to the textual pathology reports where the area of cancerous cells would be highlighted, providing a link between the textual and the image information. This is an example of fusing multimedia information snippets extracted from various sources for better answering the information needs of the users.

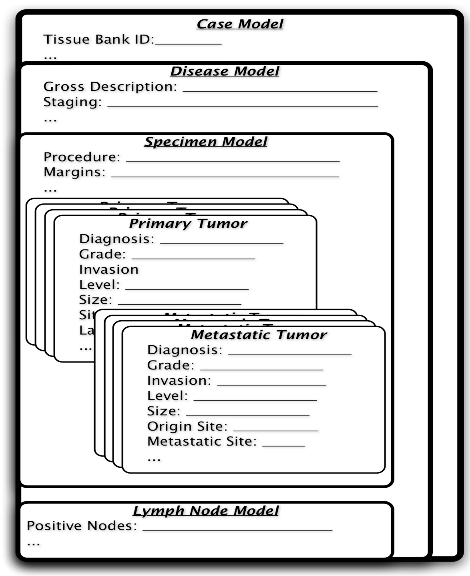

Figure 5. Case Model and its related sub-models for representing information in pathology reports

4. OPPORTUNITIES AND CHALLENGES

The realization of the concept-based health record organization and its integration in the healthcare organizations could potentially result in improvements in the process of care, decrease of medical errors and reduction in costs. In this section we like to enumerate a number of such opportunities and point out also to a number of hurdles that need to be overcome before the concept-based health records could be materialized.

Assuming that EHR becomes available some time in the near future, three main issues need to be resolved before a full fledged concept-based medical record can be created. First, is determining a common and unifying architecture such as the one shown in Figure 2.that allows for pluggable specialized analytic engines required for structuring the contents of the documents residing in the EHR. Such architecture should support standardized interface for connecting to different analytical engines developed by different entities. Second, the concepts which manifestations need to be extracted from different sources need to be defined and captured by an agreed upon ontology. UMLS or any enhanced form of it could be the start for such ontology. Third, development of robust analytic engines should be continued. For certain concepts and modalities of data that such fully automated analytic engines are yet not feasible, it would be beneficial to develop some form of semi-automated or human-in-the-loop analytics.

In the case that concept-based health records are created, one needs to investigate how the information needs of the different types of users could be mapped to the relevant portions of such records and how effective are such records in answering those information needs compared to the current approaches. One also needs to know what combination of media associated with a given concept could best provide the answer to a certain query regarding that concept. In other words what the incremental information gain of using multiple media is. Does a certain category of user with specific information need and a given task benefit from seeing an image snippet along with a text snippet or the text snippet is just enough for answering to her information needs. Such analysis is necessary in order to objectively justify the concept-based health record organization. Any integration and use of this new paradigm for health record organization in the clinical setting and its related tasks need to be seamless to the clinicians and other users.

In the rest of this section we give an overview of two of our work that illustrates some benefits of the concept-based health record.

4.1 MULTIMEDIA-GUIDED STATISTICAL DECISION SUPPORT

An important advancement in clinical decision support systems was taken by a recent work that proposed the use of multimedia analysis for enabling statistically-guided decision support [24]. The key idea in our approach is as follows. Physicians primarily make judgments based on single sample evidence provided by the given patient data and using their a priori knowledge of such data sets. In a futuristic diagnostic scenario, such diagnosis can be validated by systems that can pool together related data sets along with their corresponding diagnosis. For example, the system may return a response indicating that the physician's diagnosis agrees with 85% of other physicians who looked at similar data sets to those examined by the current physicians. Such statistically-guided decision support can enable physicians to make more informed diagnosis by combining their judgment with those of several other physicians who have looked at similar data sets. The task of determining similarity between the current data set examined by the physicians to those examined by other physicians is determined by the decision support system. Multimodal concept-based medical records can be invaluable for such a similarity analysis, as they collect all relevant medical information relating to the diagnosis under a single concept.

A multi-modal mining workbench that can seamlessly extract, analyze and correlated information from such concept-based medical records is now underway. Our goal in this work is to provide enabling information to aid decision support based on deeper multi-modal analysis while still leaving the judgment to the end users (providers, payers, patients). The current focus is on demonstrating this capability in the context of cardiovascular diseases. Specifically, we are developing automated methods of extraction of healthcare metadata from anonymized cardiac echo videos, EKG time series, and patient treatment information available from CDA and HL7 documents for storing in a concept-based medical record. We are developing algorithms that search, mine and correlate information in the metadata repository to derive intelligent reports about trends, statistical validations of diagnosis, and statistics about effectiveness of treatments.

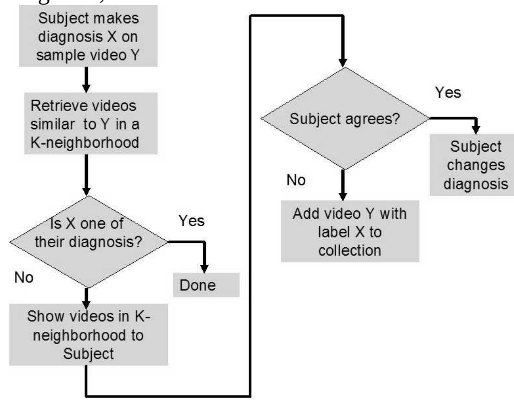

Figure 6. Flow diagram of the video guided decision support system

4.1.1 EXTRACTING DISEASE-SPECIFIC INFORMATION FROM ECHO VIDEOS

Cardiac echo videos provide information about cardiac mechanical performance which is of critical importance in understanding the etiology of heart diseases. Analyzing the motion of the heart can serve as an excellent indication for the existence of various pathological conditions such as "Coronary Stenosis". Low blood flow to a particular region of the myocardium often causes reduction in the contractile function, resulting in abnormal motion of the heart wall during contraction. Due to the quality of the video and the complexities of underlying motion captured, it is difficult for an inexperienced physician/radiologist to describe motion abnormalities in a crisp way, leading to errors in diagnosis. Thus automated methods of extraction of motion and other attributes from cardiac echo videos to reveal disease-specific differences can be invaluable to physicians.

In this work, motion information was described using average velocity curves previously proposed [25]. As the name implies, they represent the average motion per frame and are obtained by averaging the optical flow vectors between frame-wise corresponding points. Second-order motion statistics that model variations in velocity and acceleration are extracted from average velocity curves and serve as features for computing video similarity. Such measures included velocity distribution ratios and flat peak rations. Normal cardiac motion depicts even distribution around a mean velocity. Similarly, the flat peak ratio indicates the frequency of change in the direction of motion. Lack of sharp peaks indicates sluggish motion which happens when the heart's pumping ability is affected. The features extracted from videos were used to develop a video similarity measure. Then using the ideas of collaborative filtering, a new video sample already labeled with a diagnosis is projected into a feature space derived form the average velocity curve. A neighborhood of similar videos is assembled from the training set using the K-nearest neighbor algorithm. The diagnosis labels of the videos in the similarity neighborhood are used to pool together opinions of other physicians on similar data sets and are used to validate the diagnosis. This procedure is illustrated in Figure 6.

This work represents an important application of multimedia techniques in information-based medicine enabling physicians to make more informed diagnosis by validating their judgments using those of several other physicians who have looked at similar data sets. Future work will explore the use of other multimedia data such as EKG, Cardiac CT and MRI for aiding decision support.

4.2 COORDINATED MULTIMEDIA SUMMARY

One category of potential users of the electronic health records are students and people in medical training as was mentioned in the earlier chapters. The main information need of this group of users is educational. Today this kind of information is provided to them in the form of medical “teaching files”, which are an assembly of example cases illustrating a certain abnormality and its associated diagnosis and treatment plans. Assuming the concept-based health records were available today, one could imagine assembling such information to the users by retrieving the relevant information associated to the medical concepts of interest, reinforcing, explaining and supporting them by linking them to external sources of information, such as relevant sections of medical textbooks, or passages extracted from trusted journal articles and potentially other sources, and provide such aggregated information to the users to satisfy their information needs.

The “Coordinated Multimedia Summary" (CMS) [26] is a navigable multimedia summary with its user interface implemented using the “Synchronized Multimedia Integration Language” (SMIL) technology [27] that allows the browsing of a concept space by interacting with the media snippets in a set of associated heterogeneous information resources. CMS is dynamically generated in response to a query made by the user. Queries such as: “What is X?” and “What is the diagnosis for X?”, where “X” stands for an abnormality, such as mitral valve regurgitation. CMS spatially lays out a set of media snippets illustrating a single concept in the browsing interface and provides hyper-links to traverse the concept space by interacting with the content of each of the media snippets. As the result of the interaction of the user with any of media snippets through the provided hyper-links, all media snippets update themselves in a coordinated fashion to reflect the information related to the newly selected concept.

We have created a framework (Figure 7) for assembling such CMS browsers for the domain of echocardiography. The available resources in this case are the structured digital echocardiogram videos, textual summaries extracted from electronic textbooks on cardiology, and definitions of medical terms extracted from on-line medical dictionaries.

The available media snippets are as the following:

- Video Snippets: These are segments of the echo video corresponding to a view, mode, or a cycle depending on the related concept. For example, if the concept of interest is “mitral valve”, then its related echo video snippet is a “2D color Doppler” mode taken from the appropriate view that best shows the “mitral valve”. Such information is available through the expert knowledge.
- Textual Summary Snippets: Are summaries of the content of chapters or sections of textbooks retrieved via a textbook search. Such summaries are generated using statistical sentence extraction and semantic relationships. See [28] for the details on textbook summarization.
- Term Definition Snippets: Are readable definitions of concepts of interest, extracted from consumer-oriented on-line full text articles by combining shallow natural language processing with deep grammatical analysis. See [29, 30] for more details on term definition extraction approach.

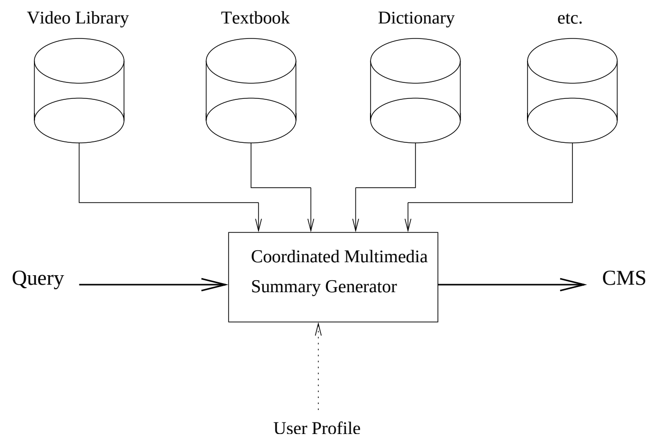

Figure 7. System for generating coordinated multimedia summaries

Figure 8 shows the snapshot of the user interface of CMS. The interface of the CMS as shown in figure consists of multiple panels. Each of those panels is associated with one of the information resources and is responsible for displaying the corresponding snippets. Users can interact with the summary by selecting the hyper-links embedded in the snippets of each panel. Note that some snippets might have multiple hyper-links embedded in them. This is because their related concept is linked with multiple other concepts in the concept space. By selecting each of the hyper-links, user initiates a move to another concept in the concept space, which translates into the update of the content of the panels with updated media snippets. The composition of the CMS interface is determined by the availability of media snippets for each concept. In the implementation of the CMS user interface, SMIL templates are being used. We have pre-designed certain number of templates that can cover the combinatorial set of media snippets that are possible for any given concept. For a selected concept, through interaction with the hyper-links in any of the panels, each resource is interrogated to find out if a media snippet is available in that resource that is the manifestation of the selected concept. Based on those responses the appropriate pre designed SMIL template is chosen and its panels are filled-out with the appropriate media snippets. Therefore, CMS is a dynamically generated multimedia summary, which at each step updates itself based on the user's information need as indicated by the selected concepts.

Such multi-source summarization and augmentation of the documents in the medical records, such as echocardiogram videos, with external sources of information such as textbooks, journal articles, dictionaries or others is only possible through concept-based access to the content of the medical records and those external sources of data.

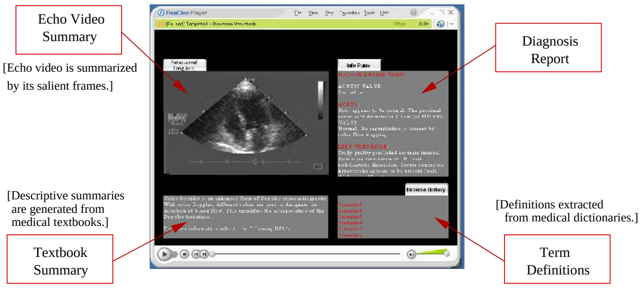

Figure 8. Interface of CMS which shows the use of multiple panels for visualizing the different forms of the manifestation of a concept

5. CONCLUSIONS

We have presented the idea of organizing health records around concepts, and shared our thoughts on what such concept-based representation involves and how could it be materialized. We also presented our opinion on where the filed of multimedia content management should be focusing in order to make it happen, and why and how the practice of medicine can benefit from such kind of health-record organization and representation. There are, of course, challenges that need to be met before one can materialize the concept-based health records. However, we believe that the field of multimedia can play a very active role in a joint effort with practitioners to take field of healthcare to the next level by facilitating the access to decision-enabling information to physicians, patients, researchers and payers.

6. ACKNOWLEDGMENTS

The first author would like to thank Chalapathy Neti and Murray Campbell for their comments and feedbacks.

7. REFERENCES

1. D. Brailer, *Key-Note Speech*. 2005: Healthcare Information and Management Systems Society (HIMSS) Annual Conference & Exhibition.
2. L.T. Kohn, J.M.Corrigan, M.S. Donaldson, *To Err is Human: Building a Safer Health System*. 1999: National Academy Press.
3. M. Sonka, J.M.F., *Handbook of Medical Imaging, Volume 2. Display and PACSMedical Image Processing and Analysis*. Vol. 2. 2002. 1250.
4. J. Jan, *Medical Image Processing, Reconstruction and Restoration: Concepts and Methods*. 2005: CRC Press. 730.
5. Y. Kim, S.C. Horii, *Handbook of Medical Imaging, Volume 3. Display and PACS*. Vol. 3. 2000. 512.
6. H.K. Huang, *PACS and Imaging Informatics: Basic Principles and Applications*. 2nd ed. 2004: Wiley-Liss.
7. ANSI., *Health Level Seven*. http://www.hl7.org/
8. NEMA, *Digital Imaging and Communications in Medicine (DICOM)*, http://medical.nema.org/
9. J.H. Kaufman, I.Eiron, G. Deen, D.A. Ford, E. Smith, S. Knoop, H. Nelken, T. Kol, Y. Mesika, K. Witting, K. Julier, C. Bennett, B. Rapp, B. Carmeli, S. Cohen, *From Regional Healthcare Information Organizations to a National Healthcare Information Infrastructure*. Perspectives in Health Information Management, 2005. **2**(10).
10. H. Feigenbaum, *Echocardiography*. 1993: Lea Febiger.
11. A.A.T. Bui, R.K.Taira, J.D.N. Dionisio, D.R. Aberle, S. El-Saden, H. Kangarloo, *Evidence-based Radiology: Requirements for Electronic Access.* Academic radiology, 2002.
12. R. Smith, *What Clinical Information do Doctors Need?* BMJ (Clinical research ed.), 1996: p. 1062-1068.
13. NLM, *Unified Medical Language System (UMLS)*, http://umls.nlm.nih.gov/
14. C. Rosse, J.L.V. Mejino, B.R. Modayur, R.M. Jakobovits, K.P. Hinshaw, J.F. Brinkley, *Motivation and Organizational Principles for Anatomical Knowledge Representation: The Digital Anatomist Symbolic Knowledge Base*. Journal of the American Medical Informatics Association, 1998. **5**(1): p. 17-40.
15. J. F. Brinkley, C.Rosse., *Imaging Informatics and the Human Brain Project: the Role of Structure, Review*. Yearbook of Medical Informatics, 2002: p. 111-128.
16. T. McInerney, D.Terzopolous., *Deformable Models in Medical Image Analysis: A Survey.* Medical Image Analysis, 1996. **1**(2): p. 91-108.
17. M.G. Linguraru, M.A.Gonzalez Ballester, N. Ayache., *Deformable Atlases for the Segmentation of Internal Brain Nuclei in Magnetic Resonance Imaging.* Acta Universitatis Cibiniensis, Technical Series, 2006.
18. S. Ebadollahi, S.-F.Chang., H. Wu. *Echocardiogram Videos: Summarization, Temporal Segmentation and Browsing*. in *IEEE International Conference on Image Processing*. 2002.
19. H. Chen, S.S. Fuller, W. Hersh, *Medical Informatics: Knowledge Management and Data Mining in Biomedicine*. 2005. 696.
20. J. Fernandez, M.Harris, C. Meyer. *Combining Algorithms in Automatic Detection of R-peaks in ECG Signals*. in *IEEE Symposium of Computer-based Medical Systems (CBMS)*. 2005.
21. S. Ebadollahi, S.F.Chang, H. Wu. *Automatic View Recognition in Echocardiogram Videos Using Parts-Based Representation*. in *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. 2004. Washington D.C.
22. K. McKeown S.F. Chang, J.C., S. Feiner, C. Friedman, L. Gravano, V. Hartizvassiloglou, S. Johnson, D. Jordan, J. Klavans, A. Kushniruk, V. Patel, S. Teufel. *PERSIVAL, A System for Personalized Search and Summarization Over Multimedia Healthcare Information*. in *IEEE/ACM Joint Conference on Digital Libraries (JCDL)*. 2001.
23. N. Kauff, M.Robson, R. Barakat, L. Norton, K. Offit, *Cancer-Reducing Benefits of Preventive Surgery may be Specific to Gene Mutation.* American Society of Clinical Oncology, 2006.
24. T. Syeda-Mahmood, D. Ponceleon, J. Yang. *Validating cardiac echo diagnosis through video similarity*. in *ACM Multimedia*. 2005.
25. T. Syeda-Mahmood, *Segmenting Actions in Velocity Curve SPace*. in *International Conference on Pattern Recognition (ICPR)*. 2002.
26. S. Ebadollahi, *Analysis of Video Content using Statistical Spatio-Temporal Models*, in *Department of Electrical Engineering*. 2005, Columbia University: New York. p. 249.
27. (W3C), *Synchronized Multimedia Integration Language (SMIL) 1.0 Specification*, http://www.w3.org/TR//REC-smil/
28. G. Whalen, *Medical Textbook Summarization and Guided Navigation using Statistical Sentence Extraction*. in *Proceedings of the American Medical Informatics Association Symposium (AMIA 2005)*. 2005.
29. S. Muresan, J.L.Klavans. *A method for Automatically Building and Evaluating Dictionary Resources*. in *Proceedings of the Language Resources and Evaluation Conference (LREC 2002)*. 2002.
30. J. Klavans, S.Muresan, *Evaluation of the DEFINDER System for Fully Automatic Glossary Construction*. in *Proceedings of the American Medical Informatics Association Symposium (AMIA 2001)*. 2001.