AALIM: Multimodal Mining for Cardiac Decision Support

T Syeda-Mahmood, F Wang, D Beymer, A Amir,
M Richmond, SN Hashmi

IBM Almaden Research Center, San Jose, CA, USA

Abstract

Diagnostic decision support is still very much an art for physicians in their practices today due to lack of quantitative tools. AALIM is a decision support system for cardiology that exploits the consensus opinions of other physicians who have looked at similar patients, to present statistical reports summarizing possible diagnoses. The key idea behind our statistical decision support system is the search for similar patients based on the underlying multimodal data. In this paper, we describe the AALIM decision support system and the underlying multimodal similarity search used for cardiac data sets.

1. Introduction

With integrated information becoming available through large patient repositories, clinicians, patients and payers are shifting from a qualitative viewing of multimodal clinical data to demanding a more quantitative analysis of information using all supporting clinical and imaging data. For example, physicians can now expect to correlate diagnostic information across patients with a similar condition to infer the effectiveness of a drug. Similarly, with large patient repositories, physicians can validate their diagnosis by seeing the agreement of their diagnosis with those of other physicians in the world who looked at similar cases. This can improve the practice of medicine allowing physicians to exploit the experience of other physicians in treating similar cases, and inferring their prognosis and the outcome of their treatments. In general, such statistically-guided decision support can allow physicians to see consensus opinions as well as differing alternatives, helping reduce the uncertainty associated with diagnosis. In the long run, this could lower diagnostic errors and improve the quality of care.

In this paper, we describe AALIM, a novel multi modal decision support system that seamlessly extracts, analyzes and correlates information from patient's electronic medical records for purposes of decision support. Since health records today have become multimodal, including images, video, text and charts, AALIM analyzes multiple modalities to identify similar patient records. AALIM then summarizes similar patient outcomes through easy-to-read reports and visualizations. Unlike previous rule-based decision support systems (see review at [2]) that are based on fixed associations of diseases and symptoms based on rule, or rule learning using neural networks [1], AALIM uses a data mining approach exploiting the vast amount of patient data and prior diagnosis recorded in a medical record system.

2. Methods

The design of the AALIM system is based on the key idea that similarity in raw patient data implies a similarity in disease labels. Thus if patients whose exam data is similar to a current patient can be found, then statistical distribution of associated disease labels can give indications of consensus in disease diagnosis. For this to be meaningful, the computation of patient data similarity has to be guided by disease-specific cues. Current work in AALIM has focused on the domain of cardiology. Sophisticated feature extraction and search techniques have been developed to find similar patients based on a disease-specific analysis of their heart sounds, EKGs, and echocardiogram videos.

2.1. AALIM system

The overall architecture of the AALIM system is described in Figure 1.

Figure 1. Illustration of AALIM system

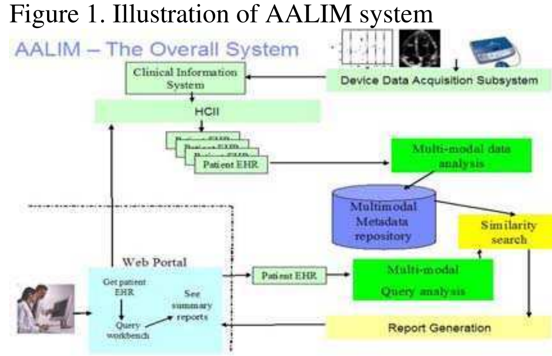

In this system, patient records from an electronic medical record system are pre-analyzed using a multimodal analysis engine. The resulting features extracted from different modalities are used to create a multimodal metadata repository along with a multimodal feature index for efficient search. The AALIM system is invoked from a browser client during a physician's examination of a patient's record. When the physician selects AALIM decision support, the current patient data is retrieved from the multimodal index and is used to

Figure 2: Illustration of search results generated from AALIM System.

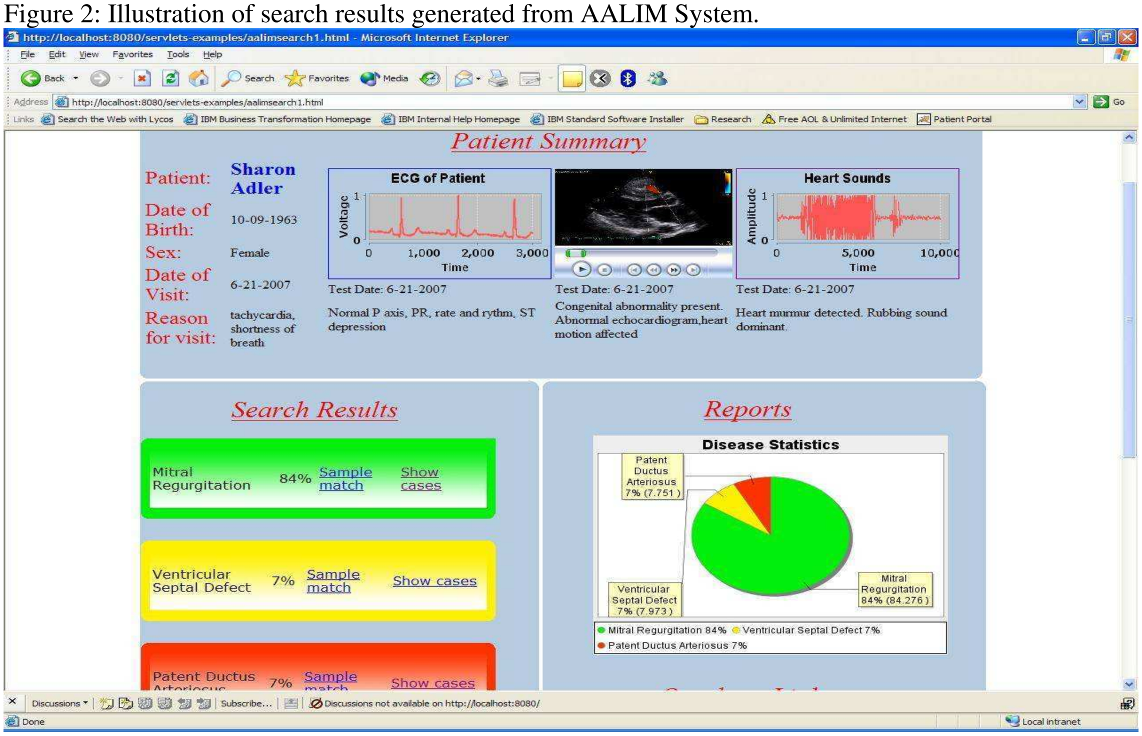

Figure 3. Illustration of identification of similar patients by AALIM system.

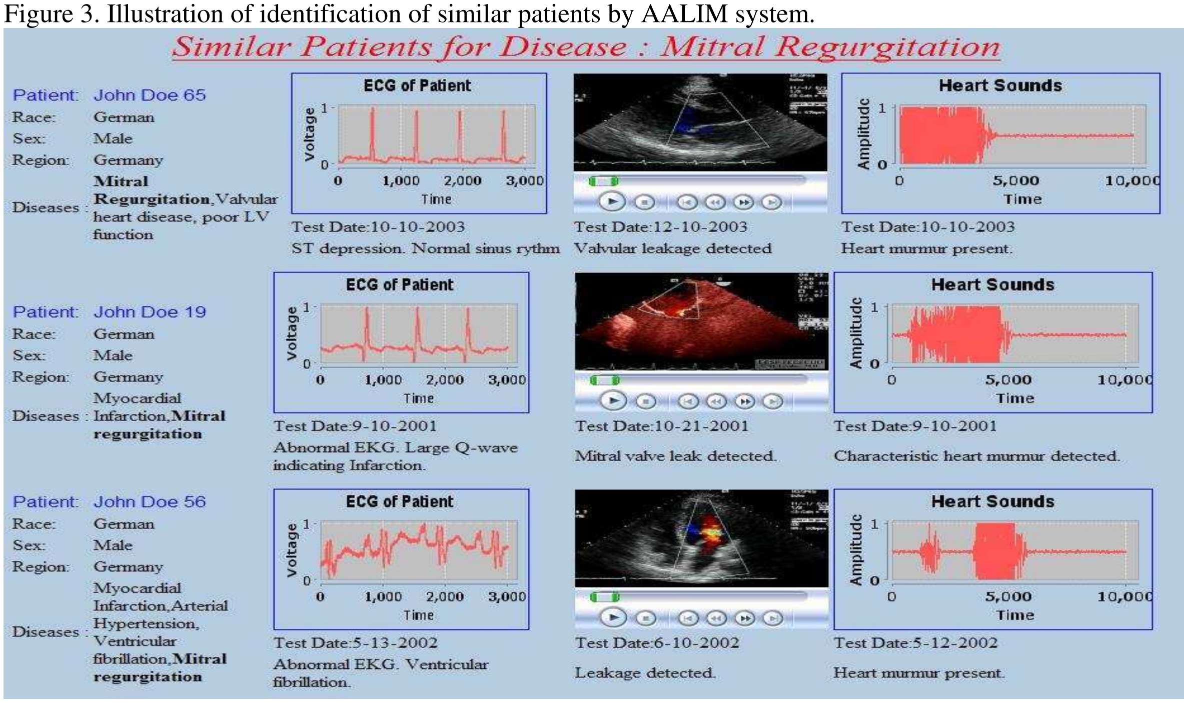

used to perform a similarity search against the rest of the patient data captured in the multimodal index. Once the set of matching patient data sets are found and their similarity ranking performed, AALIM's report generation module pools the diagnosis labels associated with the patient records to form statistical summaries of the possible diagnosis. A sample output from AALIM is shown in Figure 2.

The results of similarity search are shown grouped by diseases on the left. On the right are graphical visualizations of the consensus diagnosis based on the query patient data. For example, in Figure 2, it can be seen that the predominant diagnosis obtained from disease labels associated with similar patients is for Mitral Regurgitation.

2.2. Modality processing: ECG processing

Among the modality processing modules in AALIM are modules that process a patient's ECG, heart sounds, and echo videos. The ECGs can be derived from many sources, namely, digital ECGs, paper ECGs, and ECGs in vectorized PDF format. In each case, we process the ECG waveform to detect periodicity using a variant of normalized correlation. All channels are normalized with respect to the heart beat duration. The shape of the ECG waveform is captured as a curve with salient curvature change points recorded as features. The shape similarity of two ECG waveforms from corresponding channels is then determined using a shape matching algorithm. This algorithm is based on dynamic time warping and uses shape features derived from curvature change points from the ECG time series. Using the shape similarity algorithm, the nearest matching ECGs from the database are retrieved. The diagnosis labels of ECGs are then pooled to form a statistical diagnosis profile based on ECG.

2.3. Audio processing

In the audio processing module, heart auscultation sounds recorded from a digital stethoscope are analyzed to extract single period duration sounds. Envelopes are extracted from sound signals. Shape features are extracted from the audio envelopes and matched to other audio envelopes using the same shape matching algorithm as for ECGs.

2.4. Echo video processing

Cardiac echo videos are processed to segment various viewpoints and modes (M-mode vs. 4-chamber). 4-chamber views are processed using an active shape model [3] to obtain correspondence between different heart regions across echo videos in a database. Chambers are then extracted using a variation of multi-scale normalized graph cuts [4]. Spatio-temporal motion is extracted from each region using a deformable template registration model. Features such as septal wall motion, and ejection fraction are automatically extracted and used for disease similarity detection.

2.5. Multimodal fusion

Using ECG, echo and audio, we could get different lists of matching patients. The multimodal fusion module combines the matches from individual modules based on commonality in disease labels. Currently, we use a weighted linear combination to combine the similarity scores from individual cues.

2.6. AALIM run-time engine

The individual modality processing is carried out in an offline phase where patient data is extracted from an electronic health record system in a batch fashion. This time-consuming complex data processing is managed by a run-time workflow engine which has been designed to handle large distributed workloads. The runtime engine provides a plug-in API and the format of an XML plug-in definition file that plug-in authors use to implement their analysis component. The structure for plug-ins and the API shields plug-in authors from the details of how to access the underlying data store and the mechanics of the distributed processing model.

2.7. Cardiovascular ontology

As part of the decision support system, AALIM also provides access to a knowledge resource for cardiovascular diseases called the cardiovascular ontology. In this ontology, information about diseases is captured in the form of relationships such as causes, caused by, symptom of, risks associated with, etc. The cardiovascular ontology was generated by mining both textbook knowledge sources as well as structured knowledge sources such as Wikipedia.

2.8. Structured information processing

Structured information about the patient from electronic health records is also processed to generate statistical distributions of diseases based on demographics. In addition, this information is used to assemble reference normals for patients based on their race, gender and age.

3. Results

In this section, we present results of individual modality processing as well as the overall results produced by the AALIM system. The data sets used for experimentation came from several sources. Participating physicians from hospitals in India provided us with coherent patient data including ECG and echo videos corresponding to several diseases. Next, we used teaching collections including the Harvard data set of 301 ECGs [5], and digital ECG recordings from PTB benchmark database available from Physionet (http://www.physionet.org) which contains 12 channel 547 EKG recordings. For echo videos, we collected sample data from GE Vivid Online library, and from Yale Medical School training collection.

We illustrate the AALIM system through sample results based on the multiple modalities. Figure 4 shows the ECG of a patient diagnosed with Bundle Branch Block (BBB). The results of finding patient similarity based on ECG data is shown on the right. As can be seen, the top matches retrieved by the system are all from patients diagnosed with BBB.

Figure 4. Illustration of similarity retrieval based on ECG.

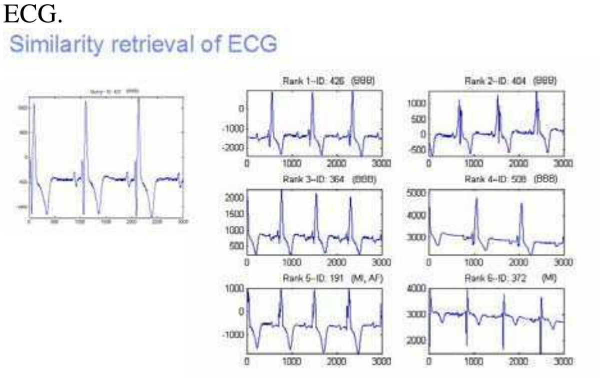

Figure 5 illustrates the results of audio-based similarity retrieval in AALIM. Here the query patient data is shown on the left, while the most similar patient based on sound analysis is shown on the right. As can be seen, the top matches based on heart sound matching correspond to patients with the same disease.

Finally, Figure 3 shows the similar patients identified by AALIM (for the query patient from Figure 2), by combining multiple cues including ECG, audio and demographics. The query patient has been diagnosed with Mitral regurgitation. The top few patients shown have all similar data as well as identical diagnosis of mitral regurgitation although some of these patients have multiple diseases.

Figure 5: Illustration of similarity retrieval based on audio.

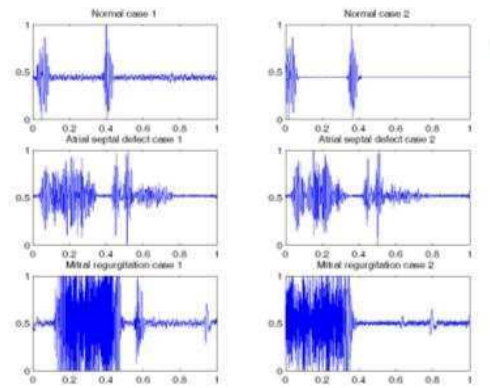

4. Discussion and conclusions

In this paper we have described AALIM, a novel multimodal search-based decision support system for cardiovascular diseases. The key idea proposed is that consensus diagnosis can be achieved by exploiting the similarity in the raw multimodal patient data. Currently, we are piloting the AALIM system among cardiologists to collect effective user study data. The results of these studies will be reported in future publications.

References

[1] Yan H. A multi-layer perceptron-based medical decision support system for heart disease diagnosis. Expert Systems with Applications 2006; 30.
[2] See review at http://www.openclinical.org/dss.html
[3] Cootes TF, et al. Active shape models, their training and applications. Comput. Vision Image Understanding 1995;61:38-59.
[4] Cour T, Benezit F, Shi J. Spectral segmentation with multi-scale graph decomposition. Proc. CVPR 2005; 1124-1131.
[5] ECG Wave-Maven, Harvard Electrocardiogram datasets. http://ecg.bidmc.harvard.edu/maven/mavenmain.asp.

Address for correspondence

Tanveer Syeda-Mahmood
IBM Almaden Research Center
650 Harry Road, San Jose, CA 95120
stf@almaden.ibm.com