

AALIM: A Cardiac Clinical Decision Support System Powered By Advanced Multi-modal Analytics

Arnon Amir^a, David Beymer^a, Julia Grace^a, Hayit Greenspan^b, Daniel Gruhl^a, Allen Hobbs^c, Kilian Pohl^a, Tanveer Syeda-Mahmood^a, Joseph Terdiman^c, Fei Wang^a

^a IBM Almaden Research Center, San Jose, CA

^b Tel-Aviv University, Isreal

^c Kaiser Permanente, Oakland, CA

Abstract

Modern Electronic Medical Record (EMR) systems often integrate large amounts of data from multiple disparate sources. To do so, EMR systems must align the data to create consistency between these sources. The data should also be presented in a manner that allows a clinician to quickly understand the complete condition and history of a patient's health. We develop the AALIM system to address these issues using advanced multimodal analytics. First, it extracts and computes multiple features and cues from the patient records and medical tests. This additional metadata facilitates more accurate alignment of the various modalities, enables consistency check and empowers a clear, concise presentation of the patient's complete health information. The system further provides a multimodal search for similar cases within the EMR system, and derives related conditions and drugs information from them. We applied our approach to cardiac data from a major medical care organization and found that it produced results with sufficient quality to assist the clinician making appropriate clinical decisions.

Keywords:

Clinical Decision Support Systems, Multi-modal Analytics

Introduction

As EMR systems are more widely adopted, clinicians are becoming deluged with data. While they may no longer need to decipher illegible handwriting in a hard-copy chart, they now have to review and understand a patient's EMR. This may include records of dozens or even hundreds of clinical encounters with multiple health care providers. As clinicians are increasingly evaluated by the quality of care they provide in diagnosing and treating patients, it becomes critical to quickly extract and understand the relevant parts of the patient's record, develop an accurate gestalt or mental model of the patient's clinical status, make a diagnosis, and then select an appropriate treatment. The speed of the analysis is especially important in Emergency Departments, where physicians make instant life-and-death decisions largely based on the patient's EMR. The daunting task of EMR systems is to support the work of clinicians by displaying a patient's record

in a complete, consistent and clear way. In this paper, we address this issue for cardiac related data of a major integrated health care organization through a system that implements advanced analytics on patient records.

One challenge in displaying a patient's health history is the fact that the EMR frequently combines records of multiple clinical encounters of different types and from different health care providers. The record of an encounter may contain progress notes, diagnoses, procedures, test results and/or any type of clinical data. Combining information from these encounters can be difficult, as the information in the records may be imprecise, contradictory, erroneous, or missing. In addition, the information might be distributed among several software systems. Physicians need to absorb this information, reconcile contradicting and/or missing data, and arrive at a diagnosis and treatment plan. To compound the problem, the treatment plan should agree with clinical guidelines, best practices developed by subject matter experts, and the clinician's anecdotal experience with patients of "similar" cases (sometimes denoted as *similar patients*).

Although most EMR systems have user friendly front-ends, they generally only provide an "encounter-by-encounter" view of a patient's longitudinal data. Reviewing a patient's record in this manner is time consuming; moreover it is also challenging to gain insight about the temporal relationships between events. One way of providing a more complete picture of the patient's history is through "problem lists" extracted from the patient's health history. Enumerating past diseases of the patient is an example of such a list. However, such lists usually hide the temporal relationships between events. More sophisticated EMR applications therefore include charts and graphs that allow clinicians to view longitudinal changes in specific tests or measurements.

This paper describes the Advanced AnaLytics for Information Management system (AALIM). The AALIM system extracts information from different data sources and presents the data in a coherent Clinical Decision Support view of an EMR (see Figure 1). It uses advanced analytics to extract diagnostic information from data such as specialist text reports as well as medical imaging. To the best of our knowledge, this is the first system to search and identify a population of patients with

similar cases, diseases, comorbidities and medications. AALIM brings the longitudinal and cross-sectional information of that population to the point of care. It provides clinicians with a consistent view of the patient's history as well as the ability to compare comparing this history with that of other patients with similar test results and disease profiles. This comparison supplies the clinician with a refined insight as to the patient's diagnosis and the comparative effectiveness of different treatments on outcomes. The system displays the information in an easy to use graphical user interface. For the initial demonstration of AALIM, we limited the scope to the field of cardiology. However, the same paradigm can be applied to other medical specialties.




Figure 1 - Simplified block diagram for the AALIM System.

Related Work

The AALIM system, originally introduced in 2007 [1], is a first-of-a-kind in that it combines clinical image data and textual information for patient-search and analysis in a clinical setting. This interdisciplinary work combines multi-modal image and video analysis, signal processing and pattern recognition, text analysis, supervised and unsupervised learning, information fusion and content-based retrieval - in which similar clinical cases are sought. In its multi-modal search for similar cardiac cases, AALIM analyzes, indexes and retrieves information from free text (cardiology reports), patient demographics, temporal data (ECG waveforms), and imaging studies (echocardiogram videos) [2,3,4]. In its goals, AALIM can affiliate itself with Case-Based Reasoning (CBR) as well as Clinical Decision Support (CDS). Furthermore, its front-end provides special visualization and usability features, designed for physician tasks. Due to space limitation we can touch upon only a few of these domains and mention related work in those areas.

Clinical decision support systems and their effectiveness have been studied for decades [5-7]. Early systems, such as MYCIN [8,9], were targeted towards explicitly capturing expert human knowledge and transferring it to other clinicians via rule-based expert systems. It is clear that in some domains rule-based systems are the ideal solution. However, they have challenges in that they are hard to develop and even more difficult to maintain given the rapid progress of modern medicine. Recent publications suggest that this issue can be addressed via case-based reasoning (CBR).

CBR produces conclusions by applying search and retrieval techniques to the EMRs of patients (for a review see [10]). Examples of CBR in the medical domain include the INRECA system [11] for diagnosing intoxications by drugs, the enhanced literature search engine developed by Coiera et al. [12], and the heart failure decision support tool by Purin et al [13,14], which searches through EMRs stored in the XML format. Most CBRs in the medical domain focus on text searches and thus ignore important information in a patients' EMR, such as medical scans and videos ([15]).

Decision support systems that analyze image information generally focus on one imaging modality. For example, Drazen et al. [16] and GE Healthcare [17] built systems specifically for ECG analysis. In [18] Patil and Kumaraswamy applied multi-modal analysis of patient data and trained a neural network to predict heart attacks. However when dealing with a large number of cardiac diseases and many combinations and comorbidities, training for each possible case is not feasible. One way to address the limitation of training for particular cases is via multimodal search systems.

Multimodal search systems fuse data from several modalities, such as text and medical scans. For example, Ruiz et al. [19] designed a multimodal search engine that combined content-based image retrieval (GIFT) with structured data (SMART) and free text (MetaMap). However, their search engine lacks automatic decision support capabilities.

In both its search engine and graphical user interface, the AALIM system uses temporal representation, integration and visualization of data. Chronological description of medical events is imperative to clinical decision making. Time line display of the patient's medical history was proposed in an early work by Cousins and Kahn [20]. Combi and Shahar [21] provided an overview of time-oriented tasks, approaches and systems in several clinical domains including cardiology,

and studied temporal reasoning extensively.

To the best of our knowledge, AALIM medical decision support system is the first multimodal search system that includes decision support. The system uses case-based reasoning [22, 1] and provides multimodal search of the EMRs of cardiac patients. AALIM requires a minimal amount of manual feedback for training as its queries are based on the current patient's status as well as similar cases. Similar cases are extracted from the database via data mining, which does not require any manual data annotation. The remainder of this article describes the system in further detail.

Methods

We previously mentioned the challenges of synchronizing information between multiple sources in an electronic medical record. In this section, we discuss our solution to transfer important information from multimodal reports into a structured diagnosis database. AALIM uses structured databases to provide clinicians with a variety of views on the state of a patient's health and on tracking of disease trends in a served population.



Figure 2 - "At a glance" overview of cardiac health.

The AALIM software package is layered on an existing EMR system deployed in a major integrated health care organization. Like many such systems, AALIM is a "work in progress", providing many features that represent a significant improvement over existing systems but still having gaps in some areas.

Rather than applying AALIM to the entire EMR at once, we focused the application on data related to cardiac diseases. We chose this strategy to target a single domain which allows us to present our system to a specific group of physicians, the cardiologists.

Data Sources

Data for this project was extracted from a variety of sources throughout the target hospital network. Depending on the source and data, the information is structured, semi- or unstructured.

Structured data acquired through routine patient care can be generally captured through fixed database tables and rigidly defined forms. In our target hospital, structured data included International Statistical Classification of Diseases and Related Health Problems (ICD9) diagnoses entered during outpatient visits, lists of significant ongoing problems, patient demographic information and medications prescribed.

In contrast, semi-structured data has defined list of fields whose content is free text. Examples of this type of data are human and computer generated data from ECG reports, the header blocks of echocardiogram reports, and the "signal" trace of an ECG report that is parsed from a printed report.




Figure 3 – Associations with comorbidities and drugs

Lastly, some data in our system comes with virtually no structure. Dictated reports for cardiac catheterizations, free text in echocardiogram reports, and streaming video from the echocardiogram procedures all fall into this category.

Once acquired, the information is stored in our data model (a relational data base plus files for binary object data) where it is made available for both analytics and presentation.

Examples of Advanced Analytics

AALIM generally first extracts features from the multi-modal data. These features are then included in the disease database, which AALIM uses to search through the data. In the remainder of this section we provide three examples for such searches.

Given a patient's ECG waveform, AALIM uses an advanced shape matching algorithm based on dynamic shape warping to find other similar ECGs. Different morphological variations of the shape corresponding to the same disease are modeled as a constrained non-rigid translation transform [2]. The ECG interpretations (a combination of computer diagnosis and cardiologist overreading) are then pooled forming a statistical diagnosis profile.

Another example is the automatic extraction of spatio-temporal motion patterns from echocardiogram videos. Our approach combines active shape model-based registration with region extraction and tracking using multi-scale normalized graph cuts [3]. These features, such as septal wall motion and ejection fraction, are added to our information about the study.

Finally, we analyze free text reports by matching keywords and phrases common in cardiac diagnosis via a taxonomy. Furthermore, we perform a shallow parse to identify modifiers and negators of these terms, for example "no evidence of mitral stenosis" or "severe case of atrial fibrillation". As can be seen, the relevant features for similarity differ from modality to modality. As new features are developed they are tested against known diseases for patients in a test set to identify how well they predict presence or absence of a disease.

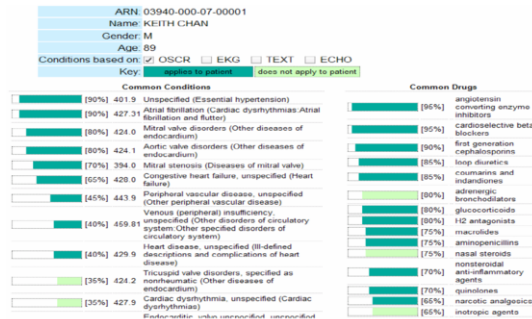


Figure 4 - "Micro cohorts" of related cases

In all cases, these additional features developed with AALIM analytics are added to the patient record as "derived" features. These features are generally considered of lower confidence than those specified by physicians. However, this type of data often provides important and valuable cues for selecting and presenting information on a patient's overall health.

Ultimately, each modality creates its own rank ordered list of most similar patients. These lists are then combined through a voting method [23] to generate a candidate list of similar patients for use in the presentation. Each modality chooses how many candidates to return based on its internal confidence thresholds.

Presentation

Our goal in the presentation of the collected information is to provide the physician with a complete overview of cardiac status of the patient. The presentation system (and thus all the underlying analytics that drive it) was motivated through interaction with clinicians, and seeks to explicitly present information about patients and diseases in a way that clinicians implicitly do when considering a patient. The difference is that instead of working with anecdotal and population studies, the system has thousands of complete medical records to draw upon for making comparisons of results and comorbidities.

We display this information via a few different "perspectives". The first perspective gives a patient overview, including a history of recent studies, current diseases and drugs (Figure 2). This gestalt view is useful especially in cases where a doctor sees a new patient that has an ongoing disease history.

The second kind of display is an overview for a selected disease (see Figure 3) such as the disease demographics, common comorbidities, and common drug prescriptions. Clinicians use this perspective to learn about common characteristics of a disease, such as the fact that mitral stenosis is more common among women than men.

The third screen (see Figure 4) helps with differential diagnosis. It shows diseases and drugs common with patients similar to the case being looked at. We define the similarity between patients by analyzing a combination of their ECG, echocardiogram, cardiac catheterization, diagnoses and other data. Users can influence this matching process by selecting specific criteria to define commonality between patients.

Heart Rate Extraction	91%
Arrhythmia detection	94.5%
Ejection Fraction	81%
Mitral Stenosis Detection	99%

Figure 5 - Selected results of the advanced analytics

Results

Today's EMR systems frequently require human intervention to accurately extract important information from multimodal reports into a structured diagnosis database. This is especially true for diagnosis information captured in non-structured modalities. An example would be a transcription of a catheterization procedure or visual evidence in an echo video.

We applied AALIM to several different collections and measured the accuracy of the tool to a given reference set. In the first experiment, we used our advanced analytics to recognize and tag patients with arrhythmia from the ECG database, which contains 20,077 12-channel ECG scans. Of these, our method could reliably extract heart rate for 18210 cases (91%). We found 1952 cases of tachycardia and 1066 cases of bradycardia. The accuracy of our approach reached 94.5% for arrhythmia patients. In all cases accuracy was computed by comparing with disease labels from an expert following the American Heart Association guidelines.

The automatic tagging capability enables AALIM to notify clinicians in the patient's "overview" screen (such as in Figure 2) of important medical conditions, such as arrhythmia. These notifications are important as they may be symptoms or complicating factors of other diseases. We note that some ECG scanners provide arrhythmia information. However, our approach can generate this tag from any (even historical) ECGs resulting in a consistent presentation of the labeling.

Ejection Fraction (EF) captures the pumping efficiency of the left ventricle, an important metric for identifying ventricular diseases. Our analytics automatically estimate the EF from the videos, which are then compared to written cardiologist reports. Out of 1604 echo videos, the algorithm agreed in 81% with at least one human manual interpretation. AALIM flags those reports where it failed to extract an EF and had to provide an estimate, and encourages the cardiologist to re-examine and correct the estimate, if necessary.

In the last experiment, we discuss the issue of diagnosis propagation. In our test set of 1173 cardiac patients, the original EMR system listed 95 patients with mitral stenosis. However, our analysis tool on the echo and textual data identified 128 cases. Of the 128 patients flagged, a cardiologist examined 105 cases in detail which included all of the 95 cases already flagged in the system. Using all available evidence, the cardiologist concluded that the assessment of disease/no disease done by AALIM was correct in all but one case, whereas the recorded diagnosis was correct only in 54 of the 95 cases.

These examples highlight the potential of advanced analytics in EMR systems. While our prototype system is still being

validated using expert supervision, the tool already allows clinician a quick presentation and a more complete view of a patient's complete health. It is clear that such a tool can be of use in many different specialties of medicine, although the underlying advanced analytics will of course differ in extracting information from the tests and procedures of each specialty.

Conclusion

Modern EMR systems are often hampered in their ability to present comprehensive, complete and coherent views of patient health due to misaligned and missing data. Such challenges can be mitigated via advanced analytics.

We found that our analytics are of sufficient quality to improve aligning records from disparate sources as well as providing estimates for missing values. We used this more complete and comprehensive data set to produce and present gestalt overviews of a patient's health status. We also identified micro-cohorts of related patients for which we computed statistics of disease prevalence and medication usage. These statistics can then assist in the differential diagnosis.

Although the AALIM system is still under development, in preliminary evaluations we found that our advanced analytics are useful to assist clinicians in improving the quality of care. Results of user studies indicate that such consolidated presentation of information is extremely helpful to clinicians, especially for emergent cases in which they may have no prior information about the patient. Further, continued development will enable clinicians to make informed decisions from multimodal analytics and promote "meaningful use" of cardiac data.

References

- [1] Syeda-Mahmood T, Wang F, Beymer D, Amir A, Richmond M, and Hashmi S N, Multimodal mining for cardiac decision support. In *Computers In Cardiology*, pp. 209-212, 2007.
- [2] Syeda-Mahmood T, Beymer D, and Wang F, Shape-based matching of ECG recordings. *IEEE Int. Conf. on Engineering in Med and Biol*, pp. 2012-2018, 2007.
- [3] Syeda-Mahmood T, Wang F, Beymer D, London M, and Reddy R, Characterizing spatio-temporal patterns for disease discrimination in cardiac echo videos. In *MICCAI*, pp. 261-269, 2007.
- [4] Syeda-Mahmood, T, Beymer D, and Amir A. Disease-specific extraction of text from cardiac echo videos for decision support, in *Intl. Conf on Document Analysis and Recognition*, 2009
- [5] Balaa Z E, Fm-ultranet: a decision support system using case-based reasoning, applied to ultra-sonography. Trondheim, Norway, June 2003.
- [6] Johnston M E, Langton K B, Haynes R B, and Mathieu A. Effects of computer-based clinical decision support systems on clinician performance and patient outcome: A critical appraisal of research. *Diagnosis and Treatment*, 120:135-142, 1994.
- [7] Coiera E, *The Guide to Health Informatics (2nd Edition)*. Arnold, London, UK, 2003.
- [8] Buchanan B and Shortli_E, A model of inexact reasoning in medicine. *Mathematical Bioscience*, 23:351-379, 1975.
- [9] Buchanan B and Shortli_E, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, 1984.
- [10] Watson I and Marir F. Case-based reasoning: A review. *The Knowledge Engineering Review*, 9(4), 1994.
- [11] Althoff K-D, Bergmann R, Wess S, Manago M, Auriol E, Larichev O, et al. Case-based reasoning for medical decision support tasks: The inreca approach. *AIIM*, 12, Hayward R S, El-Hajj M, Voth T K, and Deis K, Patterns of use of decision support tools by clinicians. *AMIA Annual Symp. Proc.*, pp. 329-33, 2006.
- [12] Coiera E, Westbrook J I and Rogers K, Clinical decision velocity is increased when meta-search Filters enhance an evidence retrieval system. *JAMIA*, 15(5), 2008.
- [13] Eccher C, Purin B, Pisanelli D, Battaglia M, Apolloni I and Forti S, Ontologies supporting continuity of care: The case of heart failure. *Comput. Biol. Med.*, 36(7-8):789-801, 2006.
- [14] Purin B, Eccher C, and Forti S, a real application of a concept-based electronic medical record. In *AMIA Annu Symp Proc.*, 2003.
- [15] Beyli E D. Advances in medical decision support systems. *Expert Systems*, 26(1):3-7, 2009.
- [16] Drazen E, Mann N, Borun R, Laks M and Bersen A, Survey of computer-assisted electrocardiography in the united states. *J. of Electrocardiology*, No. 21, Suppl:S98-104, 1988.
- [17] Hurst J W, Clinical excellence in computerized ECG analysis. white paper, General Electric, 2005.
- [18] Patil S B and Kumaraswamy Y S, Extraction of Significant patterns from Heart disease warehouses for heart attack prediction, *IJCSNS*, 9(2), Feb 2009.
- [19] Ruiz M E, Combining image features, case descriptions and UMLS concepts to improve retrieval of medical images. *AMIA Annual Symp. Proc*, pp. 674-678, 2006.
- [20] Cousins S B and Kahn M G, The visual display of temporal information. *AIIM*, 3: 341-357, 1991.
- [21] Combi C and Shahar Y, Temporal reasoning and temporal data maintenance in medicine: issues and challenges. *Comput. Biol. Med.*, 27, 353-368, 1997.
- [22] Ebadollahi S, Coden A, Tanenblatt M A, Chang S-F, Syeda-Mahmood T, and Amir A, Concept-based electronic health records: opportunities and challenges. In *Proc. of the 14th ACM int. Conf. on Multimedia*, pp. 997-1006, Santa Barbara, CA, USA, 2006.
- [23] Alba A, Bhagwan V, Grace J, Gruhl D, Haas K, Nagarajan M, Pieper J, Robson C, Sahoo N: Applications of Voting Theory to Information Mashups. *ICSC 2008*: 10-17

Address for Correspondence

Daniel Gruhl, IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120; dgruhl@almaden.ibm.com