Case-Based Teaching of Cardiac Auscultation

Kurt D. Fenstermacher
Robin D. Burke
Kristian Hammond
Department of Computer Science
University of Chicago
1100 E. 58th Street
Chicago, IL 60637
{fensterm, burke, kris}@cs.uchicago.edu

Introduction

We are developing a computer system to teach medical students and residents the art of listening to the heart for diagnostic clues, known as cardiac auscultation. The *Cardiac Auscultation Diagnosis Instruction* (CADI) system is designed to teach students both the clinical reasoning needed for effective diagnosis, and the perceptual skills needed for listening to the peculiar sounds of the human heart. Currently, learning auscultation is difficult because it requires students to visualize complex, dynamic structures as they are described by an instructor, and to focus on clinically important sounds in a sonic morass. The CADI environment will provide a realistic setting in which students "learn by doing". Furthermore, the environment will provide support for beginning students while still providing a demanding test for more advanced users. We believe that the computer's multimedia capabilities, together with intelligent tutoring principles grounded in cognitive theory, can help medical students master the challenging skill of auscultation. CADI's design draws on earlier work in case-based reasoning and goal-based scenarios, and their combination in the case-based teaching architecture.

As a human heart beats, the blood flows through the great vessels and the four chambers of the heart. In a healthy heart, the flow is regulated by one-way valves which seal off the various chambers and the heart itself. Sound is produced by the blood passing through the valves and the valves' motion. The various sounds are labeled and include the first through fourth heart sounds (S_1 through S_4), heart murmurs and various clicks and snaps. Cardiac auscultation is the act of listening to these sounds for diagnostic clues.

The Problem: Why Teaching Auscultation is Hard

A skilled auscultator can determine much about the heart, which would otherwise require expensive and possibly invasive imaging procedures to discover. Despite its value as a diagnostic tool and its ease of use, auscultation is often a weak point in the education of medical students; it is a difficult perceptual skill which is best acquired through extensive practice with an accomplished clinician. Unfortunately, many medical schools are increasing pressure on faculty to devote more time to clinical and research work which leaves less time for teaching. The teaching of cardiac auscultation has thus been neglected, despite its increasing value for the practice of medicine. We believe that the CADI environment will offer students expert guidance in a difficult skill without placing additional demands on faculty time.

The medical school at the University of Chicago uses a typical approach to teaching auscultation. During the second-year class in physical diagnosis, students are broken into small groups for hands-on training. One of the sessions is focused on listening to heart sounds, and diagnosis based on the sounds. A group of a dozen students sit around a large conference table, and each places a stethoscope on a speaker which plays back heart sounds from a selection of patients. After listening for a few moments, the students and instructor discuss what each has heard, and its clinical significance. But students and teachers are hampered by an inability to reproduce the sounds exactly, contrast different sounds or listen to sounds as they view animations. We seen have instructors become frustrated while trying to convey images in the teacher's mind to students. Often physicians will tap out the rhythm of the sound with a finger on the table or gesture with their hands to illustrate blood flow. But tabletops and hands cannot effectively show the dynamic nature of the heart.

CADI will be designed to present the same information students would receive in class, but using the computer's multimedia capabilities. Those same images in the instructor's mind will be rendered as color images on-screen, or even as three-dimensional animations, which highlight the relevant physical processes. Students will be able to listen to a sound, view its waveform and watch an animation of the underlying physical process simultaneously. Furthermore, students' learning opportunities will no longer be limited to the patients admitted to the hospital, or even by the availability of expert auscultators.

The CADI environment will consist of several components, some of which can be seen below [Fig. 1].

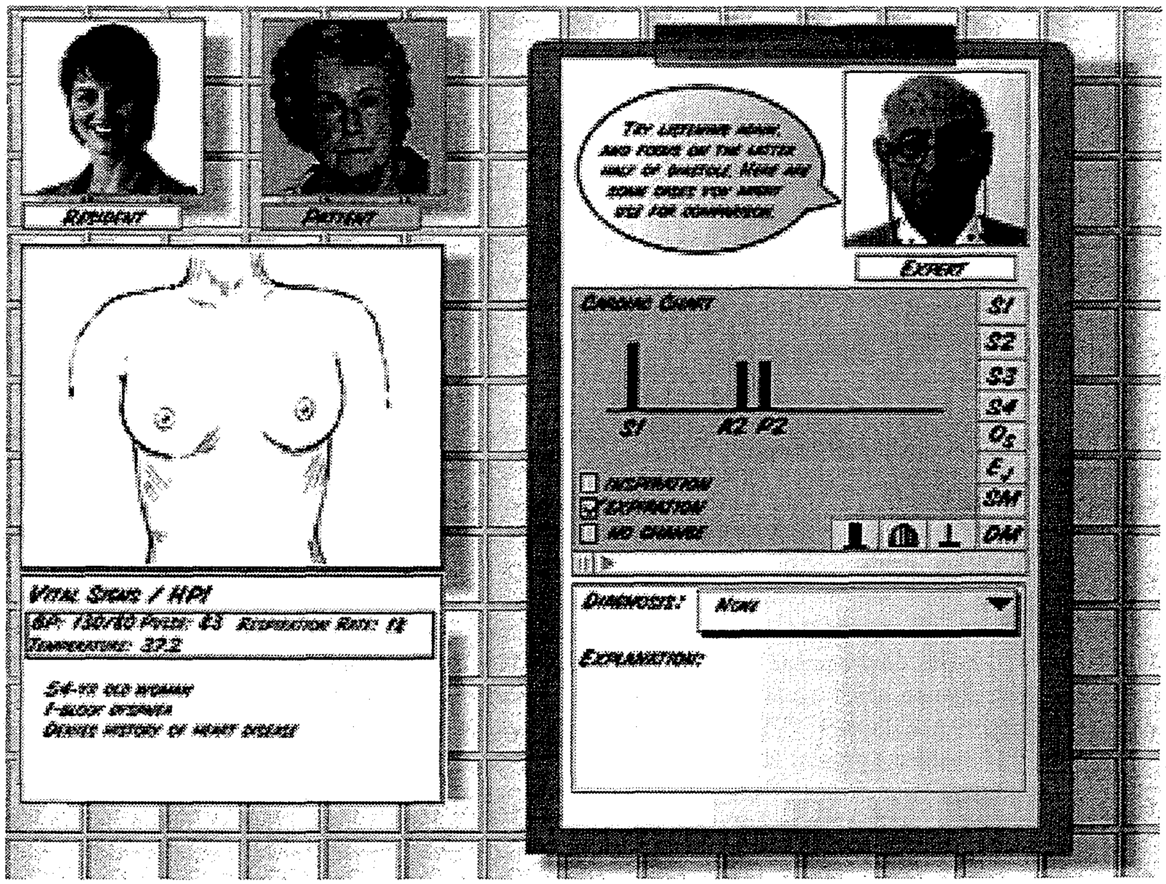

Figure 1 - The CADI Interface

After starting the program and choosing a scenario, the student will be introduced to the patient by the resident who will give a summary of the patient's condition (which will also be highlighted in a box on-screen). The student can then ask the patient to describe her problem. Next, the student would listen to the sounds of the patient's heart at the location indicated by the diagram of the chest. The right side of the figure, the CardioClipboard, is the area of the task environment where the student draws and manipulates cardiac charts, and where the student interacts with the cardiologist expert. The screen shot in Figure 1 shows a partially drawn chart. Finally, the student would use the explanation generation tools on the lower right to show the chain of reasoning which led to a diagnosis.

In building the CADI environment, we face a unique combination of challenges. First, we are trying to teach a perceptual skill – students must learn to be better listeners. But little research has been conducted regarding how people analyze auditory phenomena such as heart sounds, which lack the regularity of music or the rich context of speech. Bregman has been studying what he calls *auditory scene analysis* [Bregman 1994], which is "the process whereby all the auditory evidence that comes, over time, from a single environmental source is put together as a perceptual unit." [Bregman 1994] In CADI, the student's task is to listen to a sound complex, and then select individual components for attention. For example, a student may listen to a patient with a congenital heart defect which causes a variety of sounds which need to be analyzed individually. But the different sounds arrive at the student's ear as a jumble in which the identity of individual sounds is lost. Teaching users to analyze a mixture of sounds to separate out individual components is a task which has not been addressed by previous computer tutoring systems.

A second challenge is that students must have a firm understanding of the hemodynamic system to analyze sounds. Although the goal of CADI is to teach diagnostic reasoning, and the perceptual skill to support it, students must have a firm grasp of the complex anatomy and mechanics of the heart. A skilled auscultator's knowledge of the heart, the stethoscope, the listening location, the sounds heard and even some background knowledge in fluid dynamics all interact in reaching a diagnosis. Thus, the CADI system will need to present background information at an appropriate time and in the context of the student's current problem in a way which motivates the student.

Finally, the system must critique a student's diagnostic reasoning as he proceeds from the sounds heard to the patient's malady. There are two reasons why a student's explanation of the diagnosis is important. If a student makes an error in explaining a diagnosis, it might aid CADI in determining a student's misunderstanding of what he heard because the sounds, the diagnosis and its explanation are all closely related. Second, a detailed explanation serves as a check that the student has mastered the needed background knowledge. We approach these challenges with a design firmly grounded in theories of learning and cognition.

Theoretical Foundations

CADI's conceptual foundation rests on two theories: goal-based scenarios and case-based teaching. Goal-based scenarios address the problems of how to frame the learning task to motivate the student, and how to focus the learning task on the skills to be acquired. We use this theory to guide our choices in designing interaction with the student. We are building on the success of several previous systems that have brought together these two techniques in an overall design known as the *case-based teaching architecture* [Schank 1991].

Like other goal-based scenarios, CADI is organized around a compelling task (or mission) in a role that students will readily adopt. The task environment in CADI is a realistic role-playing simulation of hospital rounds, a typical context in which auscultation would be employed. This task provides a realistic context for the skill of auscultation and enhances the incentive to do well. While on rounds, the student plays the role of an internist who is accompanied by her residents. The residents present each patient to the user, occasionally suggesting problems they believe the patient has. The student then evaluates the patient. The simulation allows students to assume a role that they normally would not play for at least several more years. This has a natural motivating influence, but more importantly the introduction of the resident provides a natural mechanism of scaffolding the student's auscultation skills.

The student's task is to listen to sounds, and record them correctly on a chart and then explain what conditions that evidence suggests. Since this task is difficult for novices, we use the techniques of *scaffolding* and *fading* [Collins, Brown et al. 1989] to bring the goal within reach. In CADI the scaffolding for auscultation is provided by a medical resident who is presenting cases during rounds that the student is supervising. In the early scenarios, the resident will make an accurate presentation of the patient and will have perhaps made a few minor errors in the patient's chart. The student will be called upon to correct the chart so that it accurately represents the patient's heart sounds. As scenarios increase in difficulty, the resident will do less and eventually the entire chart-drawing task will fall to the student. The student will also be called upon to engage in more extended procedures, such as requiring the patient to perform physical maneuvers, auscultating at numerous points, and isolating heart sounds when lung sounds are also present.

Scaffolding for the process of diagnosis is provided by a cardiologist expert. At the earliest stages of the student's learning, the expert will provide a diagnosis and a partial explanation for how that diagnosis fits the available evidence. The student's task will be to complete the explanation, calling upon CADI's resources such as the hypermedia textbook, to help build an understanding of cardiac physiology and pathology. As the student gains in proficiency and knowledge, that task will be increasingly under the student's control to the point where the student is proposing diagnoses without prompting. The CardioClipboard will contain a mechanism for composing explanations that relate signs, including heart sounds, to diagnoses. For example, a student should be able to relate a crescendo murmur just before the first heart sound (S_1) to mitral stenosis with an explanation that describes how narrowing of the mitral valve causes the left ventricle to be filled more slowly than normal. This creates a pressure differential between the atrium and ventricle. This pressure difference is increased in atrial systole, which just precedes S_1, causing the presystolic crescendo murmur [Turner, Gold 1984]. The tools on the CardioClipboard will enable a student to construct similar explanations, with help from the cardiology expert if necessary.

Educational case retrieval

While goal-based scenarios help to structure how to frame the learning task, case-based reasoning offers insight into *what* should be taught. We believe that CBR offers a unique approach to teaching the perceptual skill of listening to heart sounds. Ideally, students should learn to hear what experienced cardiologists hear. Students must also learn to screen out sounds, and focus on a particular aspect of a sound complex in a process known as *focused listening.* We see the students' learning as a two-fold process. First, students must learn a library of prototypical heart sounds. They must become familiar with the sound pattern that a patient with, for example, an ejection click exhibits. Second, they must be able to evaluate the similarity between the sounds that they hear and those prototypes with which they are familiar. (This task is made more difficult by the nature of heart sounds, which are low frequency and often low volume.) These two aspects of students' learning are well captured by the model of human cognition known as case-based reasoning [Schank 1982; Hammond 1989]. In this model, new experiences are understood by referring to previous experiences stored in memory. To retrieve a relevant experience from memory, the reasoner must possess a similarity metric by which different cases can be compared.

Previous work in the Case-based Teaching Architecture [Burke 1993; Edelson 1993] has concentrated on the use of recalled cases to teach planning and design skills. These systems have shown that cases, particularly when presented in multimedia format, provide compelling educational experiences when integrated with a learning-by-doing environment. We plan to build on these successes to build CADI using similar principles but teaching perceptual skills.

In the Case-Based Teaching Architecture, the tutorial case presentation component of the system watches over the student's shoulder and presents relevant cases to assist in the student's learning. The expert cardiologist in CADI monitors students' actions in the task environment and responds to students in two ways. When students have made listening errors, the expert recalls examples of other patients whose heart sounds can be compared with the current patient to help improve the student's listening. When the student is in the diagnostic phase, the expert will make available parts of the hypermedia cardiology textbook that are relevant for making the diagnosis.

When the user selects a case, it become superimposed over the current patient, and many of the same functions are available. Unlike the task environment, however, the case interface provides additional tools for understanding heart sounds. Cases always have correct charts and correct diagnoses associated with them. They also have associated waveforms so that students can get direct visual correlates of the heart sounds. Students will be able to compare these features across a range of cases that the expert has chosen. For example, the expert might recall a case in which a patient has a loud opening snap and one without a snap. Through comparison of these cases and the current patient (who might have a less noticeable opening snap), the student will get a feel for what to focus on in listening for this feature.

Case Retrieval

Research in the case-based teaching architecture has shown that careful attention must be paid to the question of relevance in case retrieval [Burke, Kass 1995]. Students in an educational context do not necessarily have the ability to select truly useful examples because they are not aware of their own missing knowledge or perceptual difficulties. At the same time, it is not profitable to flood students with examples, each of which has some educational relevance. Cases must be selected carefully, with an eye to what will best enhance learning.

Educational case retrievers therefore employ pedagogical case retrieval strategies: strategies for selecting cases that take into account the specific needs of the student's learning process. Existing case-based teaching architecture systems have concentrated on planning and design tasks [Burke 1993; Edelson 1993], which are different from perceptual tasks, such as auscultation. One of the research tasks in CADI therefore will be the development and implementation of a library of retrieval strategies targeted for perceptual learning tasks.

We already have developed some ideas about the types of strategies that will be needed. Each retrieval strategy will look in the case library for cases with heart sounds that relate to the student's situation in particular ways. The following four retrieval strategies are described in terms of the kind of sounds they are intended to recall:

More distinct feature This strategy recalls a sound in which some sound feature is more noticeable than it is in the current patient. Because heart sounds and murmurs are often difficult to detect, it is often useful to hear a patient in which a particular feature is very salient in order to sensitize the ear to the more subtle version found in the current patient. For example, if the student fails to hear a soft murmur, the sound can be contrasted with a patient whose murmur is louder.

Absent feature This strategy recalls a patient whose heart sound is missing one of the features found in the sound for the current patient. The idea here is to contrast the current patient with an example that is less complex. The current patient may have a sound feature which is difficult to hear amid the complex of sounds for that patient. By recalling a patient with a similar sound, but without the complexity of the current patient, the student might then find it easier to perceive the difficult features of the current patient.

Confusable features This strategy recalls cases with sounds that are known to be frequently confused with the sounds heard in the current patient. From experience, teachers know that some features of heart sounds and murmurs are easily confused with others. For example, students often mistake S_3 for the pulmonary component of S_2, and vice-versa [Tilkian, Conover 1993]. By hearing examples of cases involving both instances, and then comparing and contrasting them, students can learn to distinguish the two.

Students' incorrect feature This strategy uses the charts that students draw as the basis for retrieval. It looks for a case that is similar to what the student has drawn rather than a case that is similar to the patient that the chart is supposed to represent. If the student creates a chart which indicates, for example, a holosystolic murmur when the murmur is confined to the latter part of systole, the system will retrieve a case for which the student-drawn chart would be correct---effectively saying "if that's what were going on, the sound would be like this."

Conclusion

Learning cardiac auscultation is difficult because it requires frequent practice, there are few experts available and it requires a deep conceptual understanding of the hemodynamic system. CADI will offer students the opportunity to practice listening on their own schedule, and from a wide range of patients. A carefully designed tutoring system, integrated with advice from acknowledged experts, will ensure that students receive the necessary support. The integration of case retrieval and comparison trains students to identify important features in heart sounds without resorting to unrealistic distortion of the sound complex. Finally, the extensive use of multimedia will give students a unique perspective, to ensure they are not only accomplished listeners, but can also reason deeply based on their knowledge of the heart.

References

[Bregman 1994] Bregman, A. S. (1994). Auditory Scene Analysis. Boston, The MIT Press.

[Burke 1993] Burke, R. (1993). Representation, Storage and Retrieval of Tutorial Stories in a Social Simulation. Institute for the Learning Sciences. Chicago, Northwestern University.

[Burke, Kass 1995] Burke, R. and A. Kass (1995). "Supporting learning through active retrieval of video stories." Journal of Expert Systems with Applications **9**(5): In press.

[Collins, et al. 1989] Collins, A., J. S. Brown, et al. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. Knowing, learning and instruction: Essays in honor of Robert Glaser. L. B. Resnick, Lawrence Earlbaum Associates.

[Edelson 1993] Edelson, D. (1993). Learning from stories: Indexing, reminding and questioning in a case-based teaching system. Evanston, The Institute for the Learning Sciences, Northwestern University.

[Hammond 1989] Hammond, K. J. (1989). Case-Based Planning: Viewing Planning as a Memory Task. Boston, Academic Press, Inc.

[Schank 1982] Schank, R. C. (1982). Dynamic Memory. Cambridge, Cambridge University Press.

[Schank 1991] Schank, R. C. (1991). Case-based teaching: Four experiences in educational software design. Evanston, The Institute for the Learning Sciences.

[Tilkian, Conover 1993] Tilkian, A. G. and M. B. Conover (1993). Understanding Heart Sounds and Murmurs. Philadelphia, W. B. Saunders.

[Turner, Gold 1984] Turner, R. W. D. and R. G. Gold (1984). Auscultation of the Heart. New York, Churchill Livingstone.