Harnessing Technology to Open the Mind: Beyond Drill and Practice for Aural Skills

Charles Lord

Instructional software for aural training has been alive and well for quite some time. Development has moved from mainframe programs on the PLATO system to microcomputers of the Apple II family and its competitors, and more recently to the Macintosh and MS-DOS machines with their greater power and speed. Software design has also developed considerably. After beginning with simple drill and practice, the approach is now much more flexible; teachers who are computer novices can modify a packaged program to their liking through built-in flexibility. Response mechanisms are now increasingly sophisticated, as ambitious programmers have constructed expert systems which interact with learners in meaningful and intelligent ways.

Objectivist Theory

A common foundation of instructional theory underlies nearly all of this work. Although little is written specifically about this objectivist model, it undergirds much current pedagogy. There is, however, a contrasting view of knowledge acquisition; in the light of this alternate theory, the rationale supporting the instructional design of computer-based aural-training software can be reexamined. One model currently in use at the University of Kentucky attempts to utilize paradigms arising from this alternate view of knowledge.

The traditional view of knowledge and instruction can be called the *objectivist* model. Knowledge, according to objectivism, is a set of facts and/or specific skills to be learned. Since these items are objective, they can be transmitted from teacher to learner accurately and cleanly; if learned well, they will be represented in memory in such a way as to be readily accessible. The purpose of instruction therefore becomes the direct transmission of this body of knowledge from the teacher to the learner. Knowledge, as represented in the brain, exists in a hierarchy or network, each schema linked to other schemata with which it is conceptually connected. If instruction can transmit the specific components of a knowledge base, the entire framework of the hierarchy or network will eventually develop. It is assumed that, once the knowledge base is in place, its application to new learning situations of varying context and complexity will be possible. In this view, the failure of a learner to apply knowledge accurately to a new situation is a sign that the components of the knowledge base are not yet in place, and that further work needs to be done on them.[1]

Problems with Objectivist Theory

What's wrong with this model for aural training? Several things, actually. The most important is that at best it works only partially. We as theory pedagogues drill and train intervals and sonorities for four semesters, and students still have difficulty identifying them accurately, much less applying their knowledge in musical contexts. Let us examine some reasons *why* this method seems to have such limited success.

1. Objectivist theory assumes the *transfer* of knowledge from one context to another. In practice, this transfer may not occur with sufficient predictability.[2] While students can memorize the sound of an isolated major sixth well enough to succeed on a theory test, they may have difficulty transferring that identification ability to actual melodies or to complete musical textures.

2. Training by drill develops cognitive schemata that tend to be rigid rather than flexible.[3] For example, learning the major sixth as $\hat{1}$-$\hat{6}$

[1]For an exposition of what is by this date traditional schema theory, see David E. Rumelhart, "Schemata: The building blocks of cognition," in *Theoretical Issues in Reading Comprehension* , ed. Rand J. Spiro, Bertram C. Bruce, and William F. Brewer (Hillsdale, NJ: Lawrence Erlbaum Associates, 1980), 33-58.

[2]Rand J. Spiro et al., "Knowledge acquisition for application: Cognitive flexibility and transfer in complex content domains," in *Executive Control Processes in Reading*, ed. Bruce K. Britton and Shawn M. Glynn (Hillsdale, NJ: Lawrence Erlbaum Associates, 1987), 177-199. See page 179 and following.

[3]See, for example, both Richard L. Coulson, Paul J. Feltovich, and Rand J. Spiro, "Foundations of a misunderstanding of the ultrastructural basis of myocardial failure: A reciprocation network of oversimplifications," *The Journal of Medicine and Philosophy* 14 (1989): 109-146; and Paul J. Feltovich, Rand J. Spiro, and Richard L. Coulson, "The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions," in *Cognitive Science in Medicine: Biomedical Modeling* , ed. David A. Evans and Vimla L. Patel (Cambridge, MA: MIT [Bradford] Press, 1989), 113-172.

of a major scale or as the opening of the NBC pattern imposes limits on the context in which the interval will be readily recognized. The learner will have to think differently to recognize the interval as $\hat{2}$-$\hat{7}$ of a major scale (which of course also has a different harmonic implication).

3. Isolating elements from their natural context emphasizes the separateness of the elements rather than their integration. If learning activates prior knowledge, teaching sonorities out of context may link them cognitively more to other isolated sonorities rather than to the role each sonority might play in a phrase. In fact, research in other domains has shown that such training can actually develop barriers between schema types instead of fostering interconnectedness.[4]

Ill-Structured vs. Well-Structured Content Domains

If, then, problems such as these exist with objectivist theory, where can answers be found? Encouraging efforts to solve similar problems in other content domains have been under way since about the middle 1980s. They span areas from reading comprehension to medical diagnosis; all these domains are what researcher Rand Spiro calls *ill-structured*.[5] Ill-structured domains share certain features. Two of those have clear applications to aural training.

1. "The same features assume different patterns of significance when placed in different contexts."[6] For example, the major sixth cited earlier carries different connotations in different tonal settings. As NBC, it may imply a tonic outline or perhaps a resolving dominant 13th chord, while as $\hat{3}$-$\hat{8}$ of a minor scale it will most likely be heard as a I^6. (Of course, if the minor key context suggests VI or a decorated III, the cognitive associations will differ still further from that which was learned.)

---

[4]See, for example, both Mark H. Burstein and Beth Adelson, "Issues for a theory of analogical learning," in *Artificial Intelligence and the Future of Testing*, ed. Roy Freedle (Hillsdale, NJ: Lawrence Erlbaum Associates, 1990), 137-172; and Carl B. Bereiter and Marlene Scardamalia, "Cognitive coping strategies and the problem of 'inert knowledge'," in *Thinking and Learning Skills: Current Research and Open Questions*, ed. Susan F. Chipman, Judith W. Segal, and Robert Glaser (Hillsdale, NJ: Lawrence Erlbaum Associates, 1985), vol 2, 65-80.

[5]Rand J. Spiro, et al., "Knowledge acquisition," 184.

[6]Ibid.

2. "An explosion of higher order interactions among many relevant features introduces aspects of case novelty,"[7] making each new case a challenge in itself. Not only can the harmonic background suggest different interpretations of an interval: so can aspects such as metric placement, rhythmic duration, articulation, accent, dynamics, and register. So many factors come into play in one's interpretation of an interval in context that its identification becomes a highly complex cognitive task.

Constructivist Theory

To deal with such complexity, educational theorists and instructional designers in other domains are turning to an alternate view of knowledge and its acquisition, commonly referred to as *constructivist*.[8] This view emphasizes a different dimension of learning from that of the objectivists. To a constructivist, the act of learning is not simply the taking in and storing of objective information for later retrieval. Learning also involves an attempt to interpret the information and to test these interpretations. Moreover, since each person's prior knowledge and experience are unique, the representation of that knowledge will differ from learner to learner. It is important, therefore, not to view knowledge as simply *received*; instead, it is the *construction* of that knowledge by the learner, the active *process* of learning, which is crucial.

That construction takes place in two ways. First, as is generally agreed by cognition experts, information learned is integrated into one's prior knowledge as it is learned, i.e., constructed into the existing schematic network. Secondly, and just as importantly, as new contexts arise in which this knowledge is to be applied, the prior knowledge itself must be reassembled or *reconstructed* to fit the new situation. Therefore, when cognitive skills need to be applied in widely differing contexts, the type of knowledge required is not so much the *retrieval* of information as its *reassembly* for flexible *application*—knowledge for *use*, not for *reproduction*.[9]

---

[7]Ibid.

[8]For a representative sampling of constructivist thought as it relates to instructional design, see *Educational Technology* 31/5 and 31/9 (May and Sept. 1991). I owe a debt of gratitude to Allen Winold for guiding me to these articles.

[9]Rand J. Spiro et al., "Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains," *Educational Technology* 31/5 (1991): 28.

As a result, the purpose of instruction shifts from the acquisition of a static retrieval scheme to the construction of knowledge for varying contexts. One means of accomplishing this purpose is to explore multiple contexts from multiple perspectives. John Bransford refers to "anchored instruction," or sustained activity within a single context.[10] One may well focus on a specific aspect within a given context for a time; the full context, however, is always present for the learner so that the interconnectedness of various aspects is apparent.

An advantage of this approach is that the most essential problem of the traditional approach—that of transfer of knowledge within varying contexts—is addressed directly. The very activities one performs in learning this way consist of working exercises within a single context, only to confirm one's success by performing similar activities in a different but similar context. If the contexts and activities are chosen wisely, a natural progression will emerge of maturing knowledge together with its application.

Implications for Instructional Design

What are the implications of these alternate paradigms for the use of technology in music instruction and the design of instructional software? To date we have seen primarily two modes of software design and hardware use: drill and practice and intelligent tutoring systems.[11] The former mode functions well for transmitting basic facts and skills, but there are limits to its applicability.

First, drill and practice designs normally do not use full contexts. In fact, rarely is any context at all provided for basic interval and sonority training. Secondly, they rely heavily on identification, i.e., the retrieval of knowledge, rather than on its construction for application. Without applying one's knowledge in varying contexts, one never learns whether or not that knowledge is confined to the specific context in which it was learned. Finally, transfer of this type of knowledge is required only within the subdomain of the program at hand, e.g., to different pitch representations of the same sonority, not to real musical contexts. The teacher never learns if the student has acquired

---

[10]John D. Bransford et al., "Anchored instruction: Why we need it and how technology can help," in *Cognition, Education, and Multimedia: exploring ideas in high technology*, ed. D. Nix and R. Spiro (Hillsdale, NJ: Lawrence Erlbaum Associates, 1990), 123.

[11]Multiple examples of both of these designs exist in the music software literature. As the focus of this paper is general and philosophical, I prefer not to cite specific examples here.

the type of knowledge that is *useful* in the larger musical world rather than merely *retrievable* in isolation. Drill-and-practice designs differ little from those of a traditional objectivist aural classroom; technology has provided an expansion of the *means* of instruction but not necessarily of its *approach.*

Intelligent tutoring designs represent a significant improvement, as they can act as guides to more accurate and fluent learning. The notion of learning through a "cognitive apprentice" is particularly attractive to constructivists.[12] Difficulties are that the approach to the subject itself is usually still objectivist, and that the production of software requires the development of an elaborate expert system to serve as cognitive apprentice from within the computer. Such systems, which are fascinating to study and create, have to be highly sophisticated to mimic human behavior, and thus often produce important insights into the underlying structure of a domain. Instructional designer David Jonassen, after building expert systems in many different domains, has concluded that the primary benefit of such systems actually comes from uncovering these insights rather than from the elusive attempt to create a fully functional system successful enough to satisfy even the unappreciative computer user.[13]

Rethinking the Function of Music Technology in Aural Training

We can look beyond these two modes of design toward a wider-ranging view of the nature and function of music technology itself, and toward the potential uses of a music workstation (by which I mean a computer, appropriate software, a MIDI keyboard, and preferably a CD-ROM drive). David Perkins of Harvard's Project Zero has identified five key facets of the various environments in which humans learn, whether these be traditional classrooms, computer labs, or other situations.[14] These facets include:

---

[12]See for discussion and examples John S. Brown, Allan Collins, and Susan E. Newman, "Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics," in *Knowing, Learning, and Instruction: Essays in honor of Robert Glaser*, ed. Lauren B. Resnick (Hillsdale, NJ: Lawrence Erlbaum Associates, 1988), 453-494.

[13]David Jonassen and Sherwood Wang, "Cognitive outcomes of building expert systems," paper presented at the conference of the Association for the Development of Computer-based Instructional Systems, St. Louis, Missouri, November 1991.

[14]David N. Perkins, "Technology meets constructivism: do they make a marriage?," *Educational Technology* 31/5 (1991): 18-19.

1. *Information banks,* or resources for learning. They include examples such as teachers in front of classes, textbooks, libraries, and computerized databases.

2. *Symbol pads,* or places where symbols are recorded and manipulated. Examples are student notebooks, blank manuscript paper, or word processors.

3. *Construction kits.* These are sets of prefabricated parts and processes which beg for creative assembly. One example is children's Lego sets, but one could also include computer-based environments in which, for example, alterations to mathematical equations produce changes in a computer-generated design.

4. *Phenomenaria.* This term, coined by Perkins, denotes spaces in which the phenomena of a domain are available for examination and manipulation. While these are more difficult to identify in educational circles, the science laboratory could qualify in one sense, while simulation games—computer-based or not—would represent another type. Also beginning to appear are computer-based learning environments, or "microworlds," such as those for the observation and manipulation of Newtonian physics.[15]

5. *Task managers.* These comprise the executive control aspects of the learning environment—managing what is to be done, providing evaluation of student progress, and perhaps offering guidance. The teacher is of course the classic task manager, but textbook exercises and most computer-based training have some of the same functions.

In light of Perkins's study of learning environments, we can examine the roles that instructional technology currently plays in music. For software designed in the drill and practice mode, the music workstation acts primarily as information bank, symbol pad, and task manager. It provides the information to be learned, serves as the interface through which the learner responds to the information, and functions as executive by managing the tasks, offering feedback, and evaluating the results. Over the last decade, the task management function has expanded to include not only programmer-determined task selection and ordering but also user-determined tasks. The soft-

[15]B. Y. White and P. Horwitz, *Thinker Tools: Enabling Children to Understand Physical Laws* (Cambridge, MA: BBN Laboratories, 1987).

ware user may determine certain aspects of the tasks chosen, their presentation modes, and the evaluation process, although this transfer of control is normally reserved for the teacher rather than the student. These programs thus allow many aspects of task management to be transferred from the instructional designer to the end user of the software, in effect providing a sort of construction kit to the teacher.

In the intelligent tutoring mode, the workstation also acts primarily as information bank, symbol pad, and task manager. What differs is largely the approach to task management. When intelligent feedback guides the learner and/or the sequence of instruction, a much more sophisticated manager plan is in operation. Moreover, the tasks themselves often require a well-developed expert system to perform the evaluation. Nonetheless, both of these design modes—drill and practice and intelligent tutoring systems—function within the same instructional paradigms. The workstation is used as an environment for designer-controlled learning, with the multiplicity of task management features actually occupying the vast majority of programming effort. Indeed, for many projects, designing the learner feedback and evaluation aspects is the most labor intensive part of the project. Unfortunately the result remains a learning experience for the student that differs little from that of the objectivist framework, with little learner control, little use of context, and little attempt to transfer learning beyond the limited context provided by the software at hand.

A Paradigm Shift?

It may be time for a paradigm shift.[16] We have at our disposal a technology that has enormous potential as a learning resource. To date we have tapped only one small fraction of that potential—the portion that derives directly from our own collective past experience. Might we not also explore the new possibilities provided by today's music workstation?

What if we thought of a music workstation less as a task manager and more as a phenomenarium—a learning environment in which all aspects of the musical world are available to study, to manipulate, and to use creatively, at least in simulated form? In this microworld one could find potentially not only pitches and rhythms but also timbres, dynamics, articulations, textures, layers, formal designs, arrangements, and notated scores. What if we allowed the learner to control

[16]Thomas S. Kuhn, *The Structure of Scientific Revolutions*, 2nd ed. (Chicago: University of Chicago Press, 1970).

much of this experience? With appropriate guidance either before or during the process we might encourage students to *explore* this world, guiding them through specific projects that focused their efforts in productive ways.

Might we not ask more open-ended musical questions, allowing students to respond in ways they themselves perceive as musical? Indeed, in a world as complex as that of music, there is usually more than one correct answer to a musical question. Why need we wait to begin such attempts until expert systems are developed to handle evaluation of such questions? Is it not possible that the technology, instead of handling evaluation for us, can be harnessed to reveal more of the cognitive processes and learning strategies that students undertake in solving musical problems? Would not these insights in fact prove even more beneficial for the development of future instructional designs?

If we did allow a shift to this approach, the result would be a vastly different learning environment for students. It would incorporate the principles of constructivist thought with the use of a music workstation as phenomenarium. It would ask questions in music and expect responses in musical sounds rather than symbols. It would allow for creativity—a feature that most students of music would relish. And it would leave evaluation mostly to the teacher, who would thereby gain greater insight into students' aural developmental processes.

An Example of the Paradigm Shift in Practice

At Kentucky, we have begun to implement such an approach. We hold at least half of our aural training sessions in our technology lab, which has eight Macintosh computers and twelve MIDI keyboards. While our efforts to date have been informal and developmental, one type of anchored instruction in a fully contextual environment has proven fruitful. We use a rudimentary MIDI sequencer, one which reads and stores MIDI input but does not provide editing capabilities. Our rationale is that we can be assured that what students have recorded is what they *played* in real time, not what they later calculated, analyzed, and edited. Into this sequencer we input a parallel period or other formal subunit from the standard musical literature, complete with its full context. We then ask students to perform various tasks with the sequence to focus on the interrelated elements of the music (see Fig. 1, left side).

Figure 1. Schematic diagram of Aural Explorations model for contextual aural training.

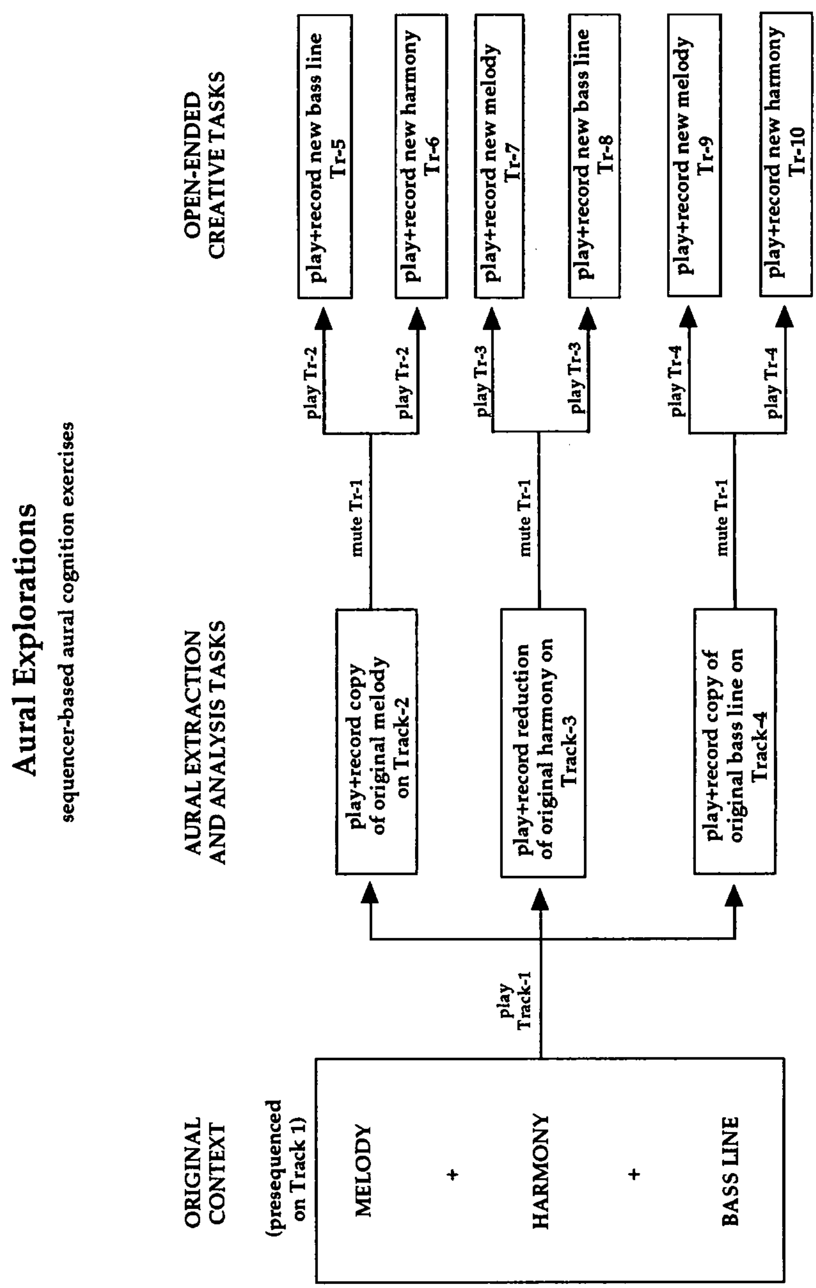

For example, one goal we have is to ensure students' ability to focus aurally on the structural framework of the passage, i.e., the outer voices and the underlying harmony. One means of training this ability is to ask the student to focus on each aspect until it has been cognitively extracted from its context. The procedure we use requires students to play back and record on the sequencer first the melody, then the essential bass line, and finally the harmony—both harmonic rhythm and sonorities—as they perceive them. Each aspect is recorded on a separate track for teacher evaluation and/or for further creative work. The results of these tasks are stored on diskette and submitted to us for evaluation.

There are many rewarding features of this process.

1. The musical context is always present for the student, thus reinforcing the interconnectedness of the various elements of music.

2. These exercises train both memory and comparison skills; comparison requires that both long-term memory and immediate perception be available simultaneously. This process simulates what students will experience later when they coach lessons or lead rehearsals.

3. By asking students to focus on each aspect separately, we ensure that the skills needed for success in other types of exercises (e.g., being able to isolate the bass line for dictation) are being trained.

4. Dual perceptual modalities are required, joining the kinesthetic with the aural. Research has shown that multidimensional or multimodal perception—what cognitive psychologists call *dual encoding*—leads to more flexible and readily applicable learning.[17] Thus training with both aural perception and the associated muscular motion produces a more cognitively flexible learning.

5. Analytical processing is required for reproducing the harmony of a passage. Reductive decisions about the number of distinct chords (harmonic rhythm), sonority identification, and, to a gradually increasing extent, voicing are all incorporated into student responses. The result is that the aurally generated analytic processes in which students engage become evident to the evaluator. Indeed, we are finding it fascinating to *hear* how students *think*.

---

[17]Cynthia Leshin, "Cognitive maps for the development of higher level thinking skills," paper presented at the conference of the Association for the Development of Computer-based Instructional Systems, St. Louis, Missouri, November 1991.

6. Much of the control of the exercise is left to the students. Clear tasks are provided, but factors such as limits of time spent, number of hearings, and number of attempts at playing back are not controlled. Students are expected to respond to this assignment as to any other in any class, or as they would in practicing for a private lesson.

7. By employing a progressive sequence of these exercises, we ensure that the ability to transfer learned skills to different contexts is being developed. This type of transfer is especially valuable since it is being performed between full musical contexts rather than between instances of isolated elements.

8. Finally, the musical experience for the students is much more like that of the rest of their musical life, situated as it is in full musical contexts and requiring a performed response rather than factual retrieval and identification.

As a complement to this extractive/reductive work, we proceed gradually toward encouraging student creativity. We may ask for any one of the extracted tracks to serve as the basis for original work (see Fig. 1, right side). For example, we may ask that the student listen only to the extracted melody, then create a new bass line to complement it. Alternatively, we might ask for a different harmonization. Similar exercises can be created for an extracted bass line or harmonic progression.

These examples constitute but a first step toward a new approach. Though we at Kentucky have only begun to explore the possibilities, several informal observations are worth noting. As teachers, we are learning more about how students hear from this type of exercise than from notated examples and are thus able to guide students more effectively. Although there is less emphasis on accurate labeling of components, students are demonstrating their ability to transfer their knowledge from context to context; the interconnectedness of the various elements becomes an inherent feature of these students' understanding of music. Finally, it is a joy for us as musicians to ask questions in music and receive responses in musical sound rather than on paper. Both students and faculty sense immediately the clearer, more direct relationship between aural training and the rest of the musical world.[18]

[18]We do not require all of the possibilities represented in Figure 1 for all sequences; rather we vary and expand them in a manner designed to derive maximum cognitive flexibility from assignments of reasonable length.

Summary

This paper has examined both the objectivist and constructivist views of knowledge acquisition and consequently of instruction. It has outlined several reasons why constructivist views deserve application to the development of computer-based aural-training software, as well as to instruction in the classroom. It also provides one example in which a constructivist approach is producing rewarding results in practice.

A constructivist vision of the aural-training software of the future might include at least these features:

- it would ask questions with musical sound and receive answers in musical sound;
- it would allow open-ended responses, i.e., multiple right answers, for complex questions;
- it would employ anchored multidimensional and multiperceptual instruction, with sustained activity in a single complex context;
- it would leave most evaluation off-line, at least for open-ended responses; and
- it would simulate the real musical world as closely as possible.

If we as instructional designers—computer-based or classroom-based—explore this approach, the music workstation will be a phenomenarium, encouraging students to learn by exploring music in new and creative ways.