

AudioExplorer: Multiple Linked Representations for Convergence

Jochen Rick

College of Computing / GVU Center
Georgia Institute of Technology
jochen.rick@cc.gatech.edu

ABSTRACT

The problem of convergence is an important one for designing collaborative-learning environments. Ideally, learning environments can allow novices to work together to achieve convergence of correct understanding, without constant support from experts. Thus, these learning environments would enable active learning by inquiry without much teacher involvement.

In this paper, I describe AudioExplorer, a learning environment where students, working in pairs, explore the physics of music. AudioExplorer is remarkable in its use of large number of multiple linked representations. I show how students worked together across the representations to achieve convergent understanding.

Keywords

Representation, software design, convergence, audio, inquiry-based learning

INTRODUCTION

Jeremy Roschelle proposes that the “crux of collaboration is the problem of convergence (Roschelle, 1996).” Can two (or more) people working together reach convergence of understanding? Furthermore, is that convergent understanding closer to real understanding than the members of the group could have reached without collaborating? Roschelle shows that students (working in pairs) can achieve convergent conceptual change, using the Envisioning Machine (EM) software. EM is a direct-manipulation graphical simulation of particle dynamics (velocity and acceleration of a particle). Students are asked to manipulate position, velocity, and acceleration of a particle to match the motion of a simulated ball. Though students did not converge on everything that scientists know about velocity and acceleration, they did manage to work together to achieve better understanding. Yet, Roschelle found that only in a few exemplary cases did the process of convergence of conception really take place.

EM provided an exploratory learning environment where two students (novices) working together achieved better understanding than they would have on their own. This process contrasts with apprenticeship, where a novice learns from an expert (Lave & Wenger, 1991). Since a typical classroom environment contains few experts (for the most part, one teacher) and many novices (students), creating learning environments where collaboration among novices is productive becomes necessary. So, what features of a learning environment are positive for supporting the process of convergent conceptual change?

In this paper, evidence is presented on how multiple linked representations can support the convergence process. In many scientific fields (such as mathematics and chemistry), phenomena can be looked at from different perspectives. For instance, in mathematics, a two variable relation can be looked at as an equation, graph, or table (Kaput, 1989). In chemistry, a chemical reaction can be looked at as the physical chemicals, the underlying reaction equations, the results of spectroscopy, etc. (Kozma et al., 2000). The power of multiple representations is that they emphasize different aspects of the same system. To understand each representation and how they are linked together is to understand the domain more completely than any one individual representation.

Both Kaput and Kozma assert that connecting multiple representations in a learning environment should be helpful for student understanding. This is based on the theory that students should be able to move between different representation and that each can inform the other. In chemistry, Kozma finds that experts move easily and often between different representations, while novices tend to get fixated on one representation (Kozma, 2001). Offering clear linkages between the representations should be a way to scaffold novices to go between them. As for convergence, multiple linked representations offer the opportunity of different ways of exploring the same domain. Since the evidence displayed by the environment is multiple (more than a single representation), there is a greater chance that useful convergence dialogue will occur.

DESIGNING EXPLORATORY TOOLS

AudioExplorer is a computer environment to explore the physics of sound by examining the frequency domain. The frequency domain is a transformation of the sound signal into its frequency components; since our ear perceives frequencies, examining the frequency domain is a useful way to understand the properties of music (Olson, 1967). The system consists of a music keyboard giving sound input into the computer (Figure 1); the AudioExplorer software displays the signal on the screen, which can then be analyzed by the students.

AudioExplorer is a tool for inquiry-based learning. The environment gives the users the opportunity to explore the subject (audio and music) and thereby discover the principles of the subject rather than passively learn about them. Thus, learning is active and students are encouraged to construct their own meaning.

AudioExplorer was designed to be used in the “Physics of Music” class at Georgia Institute of Technology. Pairs of students worked together to complete a laboratory assignment.




Figure 1: AudioExplorer usage set-up with numbers indicating representations

First, the student strikes a key on the music keyboard (point 1 on Figure 1). This produces a sound (point 2). That sound is converted by FFT (Fast Fourier Transform algorithm) to an instantaneous frequency response notation (point 3), where sonic energy is indicated by length of the line. The frequency response is recorded over time (point 4), energy being indicated by darkness. Then, the student can use the analysis line to find out the frequency of the harmonics (point 5); the students can drag the analysis line up and down by dragging on the spectrum graph to measure the exact frequencies. The calculations that the software shows (at point 5) convert that frequency to the matching key and highlight it on the display’s keyboard (point 6). In the example in Figure 1, the fundamental

harmonic is the key that was stuck originally by the student (point 1). Thus, the multiple representations come full circle.

Each representation has different features that allow the users to look at the sound phenomena from different perspectives. The physical position of the key on the keyboard (point 1) indicates the note of the sound that will be produced. It (or rather the music scale, which is roughly the keyboard layout rotated 90 degrees counter clockwise) is the usual notation for the elements of music (i.e. notes). The sound made by the instrument (point 2) is perhaps the most important representation as it is the sound that the user is looking to investigate. While we notice differences, such as changes in pitch or instrument, well in this representation, it is hard to quantify these differences. Converting the sound to the frequency response (point 3) makes it easier to analyze. Notes have energy bands at regular intervals along the frequency axis. The difference between two instruments playing the same note is in where the energy is distributed and not where the energy bands are. The distance between the intervals determines the note. For example, the same note played an octave higher will have double the interval size of the original note. Like the sound representation, the frequency response varies over time. To better analyze the sound signal over time, the frequency is recorded over time (point 4). While it is harder in this representation to tell the difference between energy levels of different instruments, temporal effects (such as fading, vibrato, and attack) are much easier to examine. Furthermore, different notes of the same instrument can be examined over time. Figure 1 shows six notes (C, C up 12 half-notes, G up 7 half-notes, C up 5 half-notes, E up 4 half-notes, and G up 3 half-notes). The last five notes' fundamental frequency is equal to the first five overtones of the original C note. The sonogram representation shows a time process (the frequency response over time) without having to rely on time; instead time is converted to the x-axis. So, whole sections of sound can be analyzed as a whole. In contrast, the frequency analysis tool (point 5) only looks at one important dimension—the exact frequency. Moving the analysis up and down let's users analyze the frequency position of each energy band. For discovering that harmonics are located at regular intervals, this representation is particularly useful. Perhaps the most difficult concept of the subject to understand is the interaction that takes place between exponential and linear relationships. A note's fundamental frequency increases exponentially: one octave up is 2 times the frequency, two octaves up is 4 times the frequency, three octaves up is 8 times the frequency, and so on. Meanwhile, the note's harmonics increases linearly: the first overtone is 2 times the frequency, the second overtone is 3 times the frequency, the third overtone is 4 times the frequency, and so on. That relationship is what the calculations in Figure 1 (point 5) are demonstrating. So, the (linear) frequency is converted to the nearest (exponential) note on the keyboard (point 6). These two tightly-linked (a change in the one immediately produces an effect on the other) representations show this relationship well.

So, each representation has different affordances that allow the user to better examine the domain. Each representation suppresses some aspects of the domain and emphasizes others, thereby supporting different forms of approaching the material. Perhaps most importantly, the linking of these representations creates “a whole that is more than the sum of its parts (Kaput, 1989).”

EVALUATING AUDIOEXPLORER

Evaluation of AudioExplorer was done for three criteria: user-interface, collaboration, and learning. Since this was the first time that the software was used with real users, a user-interface evaluation was necessary. The process of collaboration is essential to convergence, so the patterns of collaboration were evaluated. Finally, in order to demonstrate how this process of collaboration can lead to learning, the learning that occurred was evaluated.

Based on the depth of analysis needed, task observation and follow-up interview were chosen as the sources of data. Participants were recruited from the “Physics of Music” class; of the 15 students in that class, 3 volunteered, all undergraduates. Two of these, Kelly and Chris, were paired together. This left Mike without a lab partner for this experiment. So, Amber was recruited based on her interest in audio. Amber is a graduate student with some background and interest in music and audio. She played a music instrument throughout high school and has a research interest in audio; however, she has never been in a music theory class.

For the first part of the evaluation, the groups worked on the lab assignment. This took place in the a usability laboratory. The student's interaction (with each other and the software) and the screen image were captured on video. Kelly and Chris spent 56 minutes on task; after finishing the assignment, they continued playing with the system for another 21 minutes. The extra time was spent on seeing if certain musical concepts, covered in class, could be demonstrated by the system. Mike and Amber spent 46 minutes on task, deciding to end the session once the assignment was done. Based on this data, individual follow-up interviews were conducted, averaging 10 minutes.

User-Interface Evaluation

The first set of evaluation focuses on usability, because without adequate usability, learning and collaboration are impossible. Both the laboratory assignment and the software were intended to be used in certain ways. Were these used as intended? How was the interface used? Where the users able to use AudioExplorer as an exploratory tool?

Much of the interface was determined based on informal demonstrations and the intuitions of the designer; how did it work in this different environment?

Several usability problems were found with the software. First, the software did not make it easy to compare two spectrum graphs (audio signal at that time) to one another; although the data was represented on the sonogram graph (audio signal over time), it was not in a form that made for easy comparison. Neither group found the feature that clicking and dragging on the spectrum graph moved the analysis line. Both groups tried clicking on the sonogram graph, which did not do anything; instead, both groups used clicking on the computer keyboard to move the analysis line.

For Kelly and Chris, the frequency scale on the right side was not clear enough. The scale (before revision) only showed the markings for 0 and 2756 Hz. This confused them to the point of not being able to tell that it was a linear scale. They guessed that it might be exponential. Chris, between the usage and the follow-up interview, figured out that it must have been linear. Kelly still thought it was exponential by the time of the interview. This problem is largely due to the preconceptions of the designer, who never anticipated this problem; after all, if you have a strong mathematical background, you realize that an exponential scale never has a 0 on it.

The laboratory assignment was designed to introduce users to the software, elicit thinking about the subject, to guide the users through the learning tasks, and to relate the concepts illustrated to the formal language. For the last goal, a glossary was provided in the back of the booklet; during usage, neither group actually used the glossary; although, the “Physics of Music” students already used the proper musical terminology during usage. It did succeed in getting users started with the software, but the users clearly missed some of the features explained in the introduction (see above). Also, the assignment did a good job eliciting inquiry into the subject. Except for one case where Mike and Amber misread the document to say “C scale” instead of “C note,” users were able to clearly understand and interpret the questions.

Because of the nature of this task (gathering messy audio signals) and the nature of the software (simply showing the signal it received), there might be a usability problem of students misinterpreting misleading data. This possibility came up a couple of times during usage. The following is an example of where the users were able to resolve such a problem through collaboration. Figure 2 shows a sonogram graph of violin notes. The violin has a vibrato that causes the frequency bands to waver over time in a sinusoidal manner. At first, Kelly and Chris recorded a short note. By the time it stopped, the frequency band had already begun on its downward part of the sinusoid. The magnified portion of Figure 2 shows the misleading signal this created.




Figure 2: Violin Vibrato Effects

Kelly and Chris noticed this and thought it might be fading in frequency. This did not fit in with their musical knowledge. Eventually, they played the same note longer (right part of Figure 2) and were able to resolve the problem. Below is their dialog:

Kelly: That [points to screen] looks like it even decreases a little bit as it gets...

Chris: Yeah. It would be smart of them to put kind of a decay in there for a violin.

Kelly: Yeah, but it's like a frequency thing.

Chris: Yeah... Nah, I mean... ah... you know... the instrument. Of course, that's more of a thing for plucked instruments. It'll drift flat a little bit after you pluck the string. Of course, the violin, let me see. [plays note] Oh. That's what you're hearing—to give it a little expression, it goes into this

Kelly: vibrato.

Chris: Yeah.

Kelly: Let's see how that sounds. Go ahead and [motions to mouse]

Chris: Okay. [click] Here you go.

Kelly: [plays note] Now, there we go. There's our little beauty.

Chris and Kelly are able to negotiate their way across the multiple representations to solve this problem. The discrepancy shows up in the sonogram representation. By listening to the sound, they notice the vibrato. Then, by recording the note for a longer time, they confirm their vibrato explanation in the sonogram representation. The representations are linked in their understanding.

Evaluating Collaboration

Jeremy Roschelle sees collaboration as a process of convergent conceptual change (Roschelle, 1996). Since much of the learning approach for AudioExplorer is based on his theory, the next part of the evaluation concentrates on the collaboration. Was it useful? Did this process of convergent conceptual change happen? Did the users enjoy it?

As part of the follow-up interview, all participants were asked to comment on the value of working in pairs. All agreed that it was valuable. Chris guessed that “saying things out loud,” as you have to during collaboration, might have learning benefits. Amber thought that, due to the nature of the inputs (keyboard and mouse), it made it easier for two people to operate the system. Mike noted that it was more fun to work together. Perhaps the most elegant explanation came from Kelly:

Jochen: Did you like working in pairs? Do you think it helped you to understand the subject?

Kelly: Yeah. I think it did, because, you know, we could both kinda throw in our ideas about something. I mean, I had things I wanted to test out, just on my own, and then he had things too, and it was interesting to see each others... you know... own creative ways of looking at it... you know... at the end there, when we had some extra time. And, it was good, just you know, if we had a problem and the other person might understand it and be able to explain it.

Because Amber did not have the music-theory background of Mike, it made for an uneven interaction. Mike felt much more comfortable with the material and the music terminology. The other pair, Kelly and Chris, were much closer in prior understanding. This made the two interactions very different. Roschelle's observations take place with a pair of students where the background is similar (Roschelle, 1996). Kelly and Chris's interaction matched Roschelle's observations. An apprenticeship model (master / apprentice relationship) made for a better approximation of Mike and Amber's interaction than a convergent conceptual change model (peer / peer relationship) (Lave, Wenger, 1991). Part of the power of having a knowledgeable partner is that they are able to bring their prior knowledge into the inquiry as well as their focus during the interaction. Mike pointed out that Amber mainly helped him in pointing things out when the inquiry result did not match his theoretical expectations. As was demonstrated by Kelly and Chris's misinterpretation the frequency scale, having a peer-to-peer interaction does not always lead to convergence, but it does help. In an unequal interaction, such as Amber's and Mike's, it is often hard for the apprentice to correct and/or challenge the master. Kelly and Chris's dialogue below demonstrates convergent conceptual change in action:

Chris: Is that the organ right there?

Kelly: Yeah. See, I was playing this note [click], and then I was playing [plays note] this note here. I get it, this is the keyboard here. I didn't see that it was the exact keyboard. Okay.

Chris: Ah.

Kelly: See, it gives you that fifth in there.

Chris: Yeah.

Kelly: Oh. 'cause you're getting this pedal tone and that's a harmonic of that.

Chris: Hmm hmm. Plus organs have seemingly more partials than they ought to--at least the ones I've noticed. They have like the third and the fifth in there, you know, before you get to the octave. They're very soft, but...

Kelly: Yeah, but I think it's because... well, no... 'cause like I played [plays note] this note, which is right here, but for some reason it's putting that pedal tone in there so therefore the pedal tone is what's creating [Chris cuts in] the higher harmonics.

Chris: Oh. That's the octave of the pedal tone. Then, you've got the third and fifth.

Kelly: Yeah. So, it's like doing the harmonics of the pedal tone.

Chris: Yeah.

Kelly comes up with a theory of why something is happening. Chris is able to add to that his prior knowledge that organs have more partials. They are both able to converge on one conception, which is enriched by their prior knowledge. In addition, they move smoothly between multiple representations. Kelly uses the keyboard representation to investigate the sonogram representation. She realizes that the computer keyboard representation corresponds to the music keyboard. Finally, Chris is able to relate his observation that organs have more partials from the audio representation to the sonogram representation.

In contrast to Kelly and Chris's back and forth dialogue, Amber and Mike's interaction seem to be less equal. The dialogue below demonstrates Amber asking Mike for knowledge that Mike already has:

Amber: And, that is because?

Mike: The harmonic structure... like there's hardly any.

Amber: Any of these guys? [pointing to screen]

Mike: Yeah. Those are all harmonics. And so, with the organ there's none, virtually. You know. There's virtually nothing... above...

Amber: You mean the dark bands.

Amber does not understand how the "dark bands" on the sonogram representation relate to the sound representation. Mike explains how they relate. Amber may have gained some understanding from this interaction, but Mike's understanding is likely unchanged. The next dialogue features Mike undergoing conceptual change, while Amber only participates in a passive and peripheral way (agreement):

Mike: I see it. Okay. You take this here. [points to screen] You see where that line is right there?

Amber: Hmm hmm. [affirmative]

Mike: Okay. Now, when you drop it down an octave, it crosses right here. [points to screen]

Amber: I see.

Mike: Yeah. So, it's showing that when you play an octave, your harmonics overlap. If you're here [points to screen], the next octave should be right here [points to screen]... [clicks], but it's not. [inaudible, experiments around] Yeah. It works. Like if you do an octave, the next one where it crosses, the next octave, will be up from here. It's like thereish.

Amber: Yeah.

Mike moved between representations smoothly. In his explanation to Amber and his own investigations, he explores the sonogram representation using both the analysis line and keyboard representation. In contrast, Amber's earlier question about the "dark bands" is solely focused on the sonogram representation; instead of examining other representations herself to answer her questions, she simply ask her more-knowledgeable partner. In the last interaction, Mike learned from this episode, but Amber was not able to contribute to his understanding. Plus, Amber clearly didn't learn this concept in an active manner of inquiry. This is not to say that Amber was not valuable to the interaction. The below dialogue demonstrates Amber's value as she drives the course of inquiry, but it should be noted that she does not actively participate in the inquiry:

Amber: How does the uh... How did those notes come out as numbers?

Mike: Hmm. [questioning]

Amber: How'd these notes come out as numbers. I mean we know how it sounded, right? I wonder if there's a correlation.

Mike: Okay. Let's see. [clicks] 261. [clicks] 523. That's almost doub... that's doubled. Close. Yeah. The first harmonic is doubled. The second harmonic. [clicks] 783. The second harmonic is 3/2.

Amber: Does that fit?

Mike: Let's see now. It's like you take that and you double it. And, then you add 261 to that to get that. So, it's three times that. That's where your starting [click]. You double it to get here [click]. You triple it, you get here [click]. You quadruple it, you get here [click].

Amber: Okay.

Mike: I don't know if it keeps working. [does math somewhat aloud] That's what it is. If you double it [click], triple it [click], quadruple it [click], five times it [click].

Amber: Neat.

Mike: Six times it [click].

Amber: They're multiples.

Again, Mike and Amber are able to use multiple representations to better understand the domain. They had just played a sequence of notes (a sound representation). Now, they analyze the sonogram representation using the analysis line to come up with the frequencies of the notes (another representation).

In this apprenticeship relationship, it seems clear that the master (Mike) might have been better off being paired with an equal. Also, the apprentice (Amber) was not able to participate as actively in the inquiry as she would have in a peer-to-peer interaction. Still, there was value to both Mike and Amber in having each other as partners. Mike often was forced to engage multiple representation to explain concepts to Amber. Amber, on the other hand, had the advantage of having a more knowledgeable partner who could explain harder concepts.

Evaluating the Learning

The third part of the evaluation focuses on the learning. Was the software effective for learning? Were the students able to address the learning goals and understand those concepts? The “Physics of Music” students had already covered much of this content in class; was this exercise valuable to them? How did having two very different types of collaboration affect the learning?

All three “Physics of Music” students commented during the interviews that they felt that the system demonstrated the abstract concepts taught in class. As Chris pointed out, it gave you a “sense of what was going on, instead of just knowledge what was going on.” Mike remarked “to see it was cool.”

Based on the subject domain and the affordances of the software, six learning goals were chosen ahead of time. Table 1 shows the understanding each group gained during usage.

Goal	Kelly & Chris	Mike & Amber
Different instruments have different overtone patterns and that is a large part of what makes them sound different	correct	correct
Overtones are multiples of the base tone in frequency	correct	correct
Frequency increases exponentially with the musical scale	partial	incorrect
A base tone contains the major chord in its overtones	partial	correct
The first five overtones of a base tone are also notes on the scale	correct	partial
There are 12 half tones per octave because it makes a good approximation for a scale of just intonation	partial	missing

Table 1: Learning Goal Understanding (as gathered from the laboratory assignment, recorded usage, and follow-up interviews)

As can be seen from Table 1, the equal group (Kelly and Chris), fared better than the unequal group. The equal group spent more time on task and was able to get a deeper understanding of some of the concepts. For instance, in part one of the assignment, the participants were asked to give at least 5 observations, while leaving space for 7 observations; Kelly and Chris completed 6 observations before moving on and later came back to fill in a seventh; Amber and Mike, on the other hand, had a hard time coming up with 5 observations and never came back to it after finding 5.

For one of the goals, the students are asked about the relationship between the notes and their frequencies. Both had generated exponential graphs by recording a chromatic scale, playing every note from the lowest to the highest on the keyboard. Both pairs thought it might be parabolic, but Kelly and Chris were able to rely on their background knowledge to bring them to the conclusion that it was exponential.

Also, Kelly and Chris had an easier time deciding what was significant and what was misleading. For instance, both groups noticed a secondary pattern in the upper harmonics of the harmonic scale. Mike addressed this with “I don't know why it does that loop like that. It seems to be parabolic up to here [points] and that makes sense. But, why it does that? [points] It's weird.” Kelly and Chris discussed it for some time, finally deciding that it was not significant. Mike and Amber thought it might be significant, but didn't know how, so they marked it down on the answer sheet. Kelly and Chris's criteria for what was significant was higher; unless they had strong evidence, such as the phenomenon being present in multiple representations, they did not use it.

One of the fundamental learning goals was to connect the mathematics behind music to the music. For that purpose, analysis text was added to the right of the measurement bar that converted the frequency to the corresponding note and vice versa. Disappointingly, this text was never addressed by the usage; although, in the follow-up interview, the “Physics of Music” students said they understood the meaning of this text. Amber mentioned that she did not understand its meaning and would have liked to have had an explanation.

Lessons Learned

The end goal of this formative evaluation was to find out how to improve the system for future use. How could the system be improved? What learning goals should be covered in future evaluations? What research topics are of further interest?

Since this evaluation, several improvements were made to AudioExplorer. First, Kelly and Chris both mentioned wanting to compare spectrum graphs. For them, the usefulness of the spectrum was limited by not being able to compare things. Mike mentioned that he never used the spectrum graph, because he was concentrating on playing the notes and missed that it did anything interesting. As such, clicking on the sonogram graph now shows the corresponding spectrum graph at that time. Though this does not allow for comparing spectrums side-by-side, users can switch between the views fast enough so as to recognize differences. Second, the problematic frequency scale was updated so that it would be easier to understand. Though the old scale was more precise (0 and 2756 were exact values), the new scale (with markings at 0, ~1000, ~2000, and ~3000) was more understandable. Finally, on-screen balloon help was added to make it easier for users to discover the hidden features of the program, such as dragging on the spectrum graph moves the analysis line. Also, instead of just being in the introductory section, usability hints will be given throughout the laboratory assignment.

Though the students addressed all of the learning goals, their understanding was not always correct. In particular, the fundamental relationship of linear frequency and exponential notes was not fully addressed. So, in future revisions, a part of the laboratory assignment will explicitly deal with the mathematical analysis on the right side that best demonstrates this relationship.

DISCUSSION

Though this formative evaluation study has too few users to prove any claims to a reasonable degree of certainty, it does demonstrate how multiple linked representations can support convergence. Students (working in pairs) were able to use AudioExplorer to engage the subject. Roschelle's process of convergent conceptual change proved to be an accurate description for the interaction of the equal-background dyad. The other (unequal) dyad exhibited an interaction closer to that found in apprenticeships. Still, both groups moved between the multiple representations and were able to understand the links between them.

Multiple linked representations supported the convergence process. Moving between representations was prominent in both groups' dialogue as they justified their own conceptions. The representations were useful for building a whole conception that was more than the sum of its parts. Thus, the conceptions that they shared with each other were of a fairly high quality. At the same time, the representations provided an excellent way to support their positions. Finally, moving between representations gave them a reason to explicitly discuss their thought process (i.e. making their inner dialogue more accessible).

The convergence process supported the multiple linked representations. The partner's understanding was almost like another representation. When a student formed a theory based on the evident representations, their background knowledge, and reasoning, they would share it with their partner. To their partner, that interpretation of the evidence became another representation of the phenomenon. Unlike the other representations, the partner's conception is not always accurate, but it makes up for that in its flexibility and negotiability.

So, the combination of multiple linked representations and convergent dialogue proved to be mutually beneficial—their combined benefit was more than the sum of the individual parts.

Even the uneven pair managed to achieve convergent dialogue sessions as both the pupil and the master drove inquiry. This suggests that multiple linked representations can extend the benefits of convergent dialogue to a wider audience, including unevenly matched pairs that interact in an apprenticeship process. Both groups completed the

assignment and made significant progress toward deep understanding without the help of a teacher. Significant learning can be achieved with two learners and a supportive learning environment.

ACKNOWLEDGMENTS

My thanks go to the volunteers who used AudioExplorer, to Dr. Henry Valk, who gave me guidance and an opportunity to test this software on a real audience, to Amy Bruckman's CS7465 class that gave me feedback on the design, to John Maloney and the Squeak community for their programming support, and to many who have given me insightful feedback on AudioExplorer.

REFERENCES

- Kaput, J.J. (1989) Linking representations in the symbolic systems of algebra. *Research agenda for mathematics education: Research issues in the learning and teaching of algebra*, edited by S. Wagner & C. Kieran, pp. 167-194.
- Kozma, R. (2001) Material and social affordances of multiple representations for science understanding. AERA2001.
- Kozma, R., Chin, E., Russell, J., & Marx, N. (2000). The roles of representations and tools in the chemistry laboratory and their implications for chemistry learning. *Journal of the Learning Sciences*, 9(2), 105-143.
- Lave, J., and Wenger, E. (1991) *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Olson, H. F. (1967) *Music, physics and engineering*. New York, New York.
- Roschelle, J. (1996) Learning by collaborating: Convergent conceptual change. *CSCL: Theory and practice of an emerging paradigm*, edited by T. Koschmann, Chapter 9, pp. 209-248.