INVESTIGATING SOUND INTENSITY GRADIENTS AS FEEDBACK FOR EMBODIED LEARNING

Milena Droumeva

School for Interactive Arts and Technologies
Simon Fraser University
mvdroume@sfu.ca

Suzanne de Castell

Professor
Faculty of Education
Simon Fraser University
decaste@sfu.ca

Ron Wakkary

Associate Professor
School for Interactive Arts and Technologies
Simon Fraser University
rwakkary@sfu.ca

ABSTRACT

This paper explores an intensity-based approach to sound feedback in systems for embodied learning. We describe a theoretical framework, design guidelines, and the implementation of and results from an informant workshop. The specific context of embodied activity is considered in light of the challenges of designing meaningful sound feedback, and a design approach is shown to be a generative way of uncovering significant sound design patterns. The exploratory workshop offers preliminary directions and design guidelines for using intensity-based ambient sound display in interactive learning environments. The value of this research is in its contribution towards the development of a cohesive and ecologically valid model for using audio feedback in systems, which can guide embodied interaction. The approach presented here suggests ways that multi-modal auditory feedback can support interactive collaborative learning and problem solving.

[Keywords: Embodied learning, auditory display design, participatory design, sound intensity]

1. INTRODUCTION

Sound is an important part of many educational technologies, learning tools, and interfaces. It is especially vital in the design of systems where rich interaction that is contextual and embodied is the goal. In such systems, sound takes on a more prominent role in communicating and supporting activity, compared with traditional computer-mediated communication. For example, information-rich sound feedback approaches have already been applied in educational interfaces to teach mathematics and geography [1][2], as well as in interactive toys and games for pre-school aged children [3]. More recently, hands-free and eyes-free system feedback is increasingly important [6][7] with novel learning technologies such as tangible devices, mobile applications, ubiquitous computing, and multi-modal physical environments [4][5]. Auditory display standards from computer-mediated communication and data sonification have gradually made their way into mobile and ubiquitous applications.. In a multi-modal system, how can sound have a more active, communicative role that reliably and consistently conveys information and dynamically reflects changes in the interaction? Our proposed theoretical framework focuses on the development and exploration of a novel type of sound feedback that we characterize as an *intensity-based gradient approach* to sound.

2. BACKGROUND AND MOTIVATION

The overarching goal of this research is to explore and build a theoretically informed and empirically grounded model for using sound as a reliable and consistent feedback mechanism in embodied learning environments. In such exploration, it is important to examine not only the specific characteristics of embodiment, learning-by-doing, and experiential cognition in light of their relationship to feedback, but also to identify the methods and approaches to exploration that are most effective in tackling *situated* design issues.

2.1. Embodied Learning: Research and Design Challenges

Embodied learning employs perception, cognition and actions that are qualitatively different from traditional print or digital applications [8]. Contextual and physical activities require the support of an equally dynamic, cohesive auditory display system. Ubiquitous computing, ambient intelligence, and responsive environments exemplify some of the new application domains for learning through embodied interaction. We feel that responsive environments – ambient intelligent spaces - create distinctively different loci for technological activity and human engagement and require novel kinds of system design. In the case of sound, responsive environments offer different affordances for perception, cognition and engagement with sonic feedback for the user, and thus require a new investigation into all of these areas. Further, such investigation must necessarily be conducted using methods sensitive to the contextual nature of this type of auditory feedback. For this reason, we offer a case study investigating intensity-based auditory feedback using exploration methods adapted from design-based research.

2.2. Background of Auditory Displays for Embodied Interaction and Learning

There are a wide array of uses for sound in computerized, tangible and virtual interfaces depending on domain, application and objectives. Many such systems utilize confirmatory auditory feedback yet, new embodied and experientially oriented systems rely on sound displays for more complex feedback as a replacement for traditional graphical user interfaces.

For examples, non-visual interfaces and systems developed for users with visual impairments require sound to take the place of an alternative modality, representing elements normally experienced through vision [2][3][10]. AudioMath, for example, uses sonic feedback to represent basic mathematics concepts to blind children [1]. Children are asked to identify numbers and perform mathematical actions based on short-term auditory memory. While the interface is information-rich, the sound feedback does not guide children to the right answer, but only confirms when they arrive at it.

The creators of BAT: the Blind Audio Tactile Mapping System [2] have taken the auditory interface route further by creating a rich, narrative, audio interface for exploration of geographical maps by visually impaired adult users. Movement triggers the system to play local soundmarks and recorded environmental sounds in conjunction with abstract auditory icons that signify major cities, distances between locations and other contextual information. Still many sonic interfaces limit auditory information to single-sound, confirmatory-feedback displays and do not result in the kind of full-bodied, rich soundscape that users might find immersive, realistic, and informative.

3. AUDITORY DISPLAY FRAMEWORK

Two frameworks for sound design play a part in the intensity-based auditory model we propose. One is the acoustic communication model developed by Schafer and Truax [11][12]. The other draws largely on previous work in the field of data sonification [13][14]. Research in both acoustic communication and sonification, along with psychoacoustic and auditory perception studies, provide complementary guidelines for cognitive, conceptual and perceptual mapping of information to sound that are relevant to our model.

3.1. Lessons from Acoustic Communication

The acoustic communication framework suggests that in order to design ecological systems for sonic feedback, we could look for clues in the natural sound environment. This model suggests that there is a constant three-way relationship between listener, sound and acoustic environment at the centre of auditory perception and action. In the physical world, sound is constant and ambient, and we have to dynamically negotiate our attention toward it and our interpretation of it. Soundscapes are made up of many sounds in interplay with each other. These include ambient sounds that are present most of the time, sound signals that summon active attention, and soundmarks, which characterize distinctive acoustic spaces [11]. All of these elements together contribute to an environment's acoustic information ecology. They all convey different information and provide us with different cues that we put together in order to make decisions, perform actions and respond to our environment.

Furthermore, the acoustic communication framework is the basis of soundscape composition, which uses sampled environmental sounds to recreate, through audio manipulation, rich, immersive auditory environments. It seems appropriate that sound feedback in designed spaces for embodied learning emulate the models of acoustic communication and soundscape composition [11], as such environments bring the user experience closer to the real world. This is where the main challenge lies – how do we design auditory feedback that functions in a designed interactive environment, in the same way that sound functions in our surrounding natural world? One possibility put forward in this paper involves thinking of ambient sonic feedback as having a *dynamic intensity gradient.*

3.2. The Sonification Feedback Model

Sonification is a way to represent data using a continuous stream of sound driven by changes in values that results in an audible difference in the sound. It is used in environments where large information sets need to be analyzed hands-free or vision-free [14]. Auditory displays, and data sonification in particular, often function within contexts of embodied activity and/or distributed cognition. Yet issues of information and acoustic ecology, as well as cognitive implications of embodiment, lag far behind auditory perception research of pure and complex tones, spatialization and virtual 3D sound, among others Most sonification systems use synthesized tones, harmonic series, or MIDI notes to signify complex changes in data states, though there are contemporary exceptions [15][16][17]. Few systems employ realistic, environmental sounds to sonify semantically or contextually related processes or employ sound to support meaningful narrative or immersion, beyond the conveyance of information [18]. Discreet musical tones may function better for work-based situations insofar as they minimize ambiguity and allow for greater accuracy of perception and interpretation. However, novel ambient intelligent spaces for embodied learning rely on intuitive, immersive, ecological mappings between sound display and activity. They may even *require* ambiguity of feedback in order to stimulate certain types of activity and interactions. It is this paper's contention that rich everyday, "analogue" sound lends itself well to directive ambient feedback in such contexts, and can be used particularly effectively with the intensity-based gradient sound model here proposed.

One example of an ecological system from the auditory displays field is Gaver's ArKola application, which uses sampled representational (everyday) sounds, instead of musical tones [19]. The goal of this system is to create a dynamic aural representation of a complex workflow of a bottle plant that could aid users in maintaining it and making informed decisions based on the auditory display. Yet, even though ArKola fulfils its function of providing granular sonic information, it does not *guide* the user towards a desirable direction, it only 'provides the facts,' upon which one could make decisions.

Directive feedback examples are found in the fields of data sonification and auditory graphs research. Two design aspects – context (of activity, level of embodiment and type of acoustic soundscape), and perceivability (a measure of effectiveness of feedback with regard to the necessary cognitive task and action) provide the main requirements for designing ambient directive feedback. The question we ask here is, can these parameters be useful in situations where the design objective for sound feedback is immersion, learning-by-doing, physical interactivity and socialization, rather than work-based task accuracy. We suggest that sonification has in fact a great deal to contribute to understanding optimal uses of sonic feedback for embodied learning environments, as long as it's application remains sensitive to the specific characteristics and requirements of such systems. The next question is, of course, what is the most effective way to research and design such feedback?

3.3. Intensity-Based Sound Feedback

Enriching the kind of experiences that are possible for users of embodied, interactive learning environments requires a feedback system that supports the acoustic and information ecology of a physical space and provides directive feedback to the user.

Analogous to the popular children's game of "hot" and "cold" (one child looks for a hidden object while the other uses variable temperature descriptions to signify distance to the desired location), the idea of a gradient intensity of feedback seems promising. In this model, sound would respond in a subtle but helpful way to direct user actions by intensifying or de-intensifying soundscapes. This model could be mapped to any play or learning task-based activity in which users make incremental progress. Building upon our previous work, which we discuss below, the concept of sound intensity gradients has proven very useful in describing, examining and designing directive audio feedback.

While intensity-based sound feedback is a normally unnoticed part of daily life (e.g. paging your cordless phone at home and going from room to room listening to its sound intensifying), there are few studies of these everyday phenomena, and fewer still of their possible translation into design guidelines for sound feedback. Furthermore, few studies focus on perception of complex everyday *changing* sound, while taking into account context and purpose of activity, level of embodiment or familiarity and associations with the sound. A methodological investigation into all these different components of the model is needed in order to understand more fully how to better use and design such auditory direction.

The utility of this approach extends more generally to the design of educational environments where learning through doing is supported by a system of multi-modal displays. If the intensity-based sound feedback were interpreted correctly, learners would know not only *if* they are on a right track, but also, *how close* they are to completing a task or realizing a learning goal.

3.4. Auditory Design Guidelines

Research in sonification suggests that dynamic, ambient sonic displays require complex understanding of perceivability with regard to cognitive load, memory, type of activity, number of auditory sources, context and space of the interaction, as well as complexity of auditory content used in the sonification. Sonification provides us with a design framework for developing sound feedback for use in embodied learning environments, which identifies several major design elements, including data-to-parameter mapping, scaling, polarity and spatialization.

Data-to-parameter mapping refers to the choice of which data parameter is mapped to which sound variable. For example, we could map temperature to pitch, or to tempo. We could represent volume with timbre, or with amplitude. These design decisions should attempt to balance conceptual and perceptual associations of data and sound parameters. *Scaling* refers to the minimum and maximum value that a sound parameter will gradate between, driven by incoming data. This is also a significant decision. Even though humans can perceive fractal relationships between harmonic tones (i.e., we can discern that one tone is an approximate amount higher than another), there isn't an inherent sense of a scale in any particular sound. Scaling is tied to a particular design situation and varies with the context and purpose of an activity, as well as with the type of sound.

Polarity refers to the direction of gradient of change mapped between data variable and sound parameter. An example of positive polarity is when an increase in temperature is mapped to an increase in pitch. An example of negative polarity is when an increase in volume is mapped to a decrease in tempo. Decisions about polarity are important. Non-intuitive mappings may confuse users and result in inaccurate comprehension of information. Positive polarity is considered to be more intuitive than negative polarity [20][21]. In addition, the acoustic communication framework provides a lens through which we can think about the role of sonic feedback in multi-modal responsive environments for learning. This includes the concepts of acoustic ecology, or balance of sonic elements – preclusion of sound masking and unwanted diffusion or reflection; the concept of keynotes, signals and soundmarks as core elements of soundscapes; and the concept of listening positions, such as background, foreground or analytical listening. Finally, our auditory display framework rests on some fundamental ideas of psychoacoustics related to pitch, amplitude and timbre perception, jnds (just noticeable differences) in continuous sound, stream segregation and others [20][21].

4. RESEARCH DESIGN

The main questions of interest in our current study are: is intensity-based sound feedback an intuitive way for communicating information in learning or other problem-solving situations? Is it indeed well suited for experiential learning situations? Can users correctly and consistently judge intensity change, and how do different approaches to intensity (such as tempo-based, pitch-based or amplitude-based sonic changes) affect perceivability and effective interpretation of the feedback?

The main challenge in investigating intensity-based auditory feedback lies in isolating variables that influence the way participants perceive sound in the performance of a given task. Based upon the guidelines for sonification already noted above, as well as upon the above-mentioned theories/models of acoustic communication, there are two main points of interest in the investigation. The first is in the type of sound used to represent or communicate information. Recognizable sounds carry preexisting associations [11][22] and are easier to identify than unfamiliar, abstract sounds. Sounds perceived as annoying may affect perceivability in one way, while sounds that are deemed pleasant may be influential in another. Further, dimensions of sound such as timbre (quality of sound), pitch, amplitude (volume) and envelope (temporal signature) may also have different effects on users' perception [21].

The second major point of investigation is type of change representing sound intensity – that is, the data-to-parameter mapping. There are a number of characteristics of sound that could be dynamically varied, while still preserving a core quality of a sound or soundscape. Pitch, amplitude, timbre, envelope and rhythm are some of them, and, to complicate the issue, they can often be manipulated in more than one way, or have a compound effect. For example, timbre is affected both by changes in pitch, changes in envelope and effects such as filtering. Varying the speed of playback of a sound could change pitch, however, filtering of sound (attenuating or de-attenuating certain frequencies) also affects the perceived pitch, as well as the perceived amplitude [20]. The best way to fully isolate pitch from timbre and from amplitude is to only use single sine tones for investigation. However, this would preclude harnessing the rich and potentially more effective everyday sounds for feedback in experiential, embodied situations.

4.1. Design-Based Research: Issues and Motivation

One of the past research projects that informs the current investigation is a large-scale responsive environment project named socio-ec(h)o. Socio-ec(h)o is a six-level puzzle game, played physically by four players in a shared space. Its puzzles are physical configurations in space that players must achieve as a group. Immersive light and sound help players determine how close they are to solving the puzzle, guiding them towards the right configuration [23]. We termed this audio-visual feedback directive. This type of feedback seemed promising as a mechanism for supporting experiential learning (learning through exploration and experimentation), perhaps better than traditional desktop audio feedback designs, as it is ambient, dynamically responsive and hands/eyes-free.

There are few well-investigated guidelines for designing useful and informative feedback for responsive environments. The sound design approach in socio-ec(h)o was a ground-up, iterative approach. It was constructed after numerous low-tech and mid-tech participatory user workshops, where different potentials of sound regarding its immersive, ecological and narrative qualities were explored using a design-based research approach [24]. Design-based research is a situated, user-centered framework for conducting inquiry, where issues are investigated in-context, through the use of a designed system or artifact [25][26][27][28]. This approach has emerged out of a need to investigate not simply user reactions to presented stimuli, but as well, to better understand user interactions with real objects and environments, especially interactions that are not pre-determined by the designers. One of the most powerful and widely used inquiry techniques in this field are participatory and informant design workshops [26][27]. Other approaches include design games, low, mid and hi-tech prototype testing and scenario-based design.

Socio-ec(h)o exemplifies a design-based research approach, in that it is through its finished system that we are able to investigate the effectiveness, utility and implications of its components. However, by the time the socio-ec(h)o system was completed, the sound feedback incorporated so many intertwined approaches to change and types of soundscapes in addition to an equally complex theatrical lighting feedback system, that it was virtually impossible to separate out and examine the core issues of interest with respect to auditory display design, specifically.

For this reason, it was decided that undertaking several smaller studies using a design-based research approach to inquiry could more manageably explore the specific requirements of a model for intensity-based sound feedback, yet still preserve the situated user-centered approach to investigation. As one of the pioneers of using participatory workshops in design research, Liz Sanders emphasizes the importance of “make-say-do” in workshop-based research – not only listening to what users have to say, but accessing their internalized “know-how” by watching what they make, and how they do it [27], as well as, in a lesser sense, what they say and how they say it. By borrowing from this approach, low-tech workshops, which are open-ended, yet centered on specific sound feedback constraints, might, we hoped, afford a glimpse into participants’ tacit ideas and instincts about sonic feedback.

Two questions to address in the low-tech workshop presented here are: how well do sound intensity gradients work to help guide users to a goal; and do different types of sounds and different types of sound changes (varying pitch versus varying amplitude) influence how effective the feedback is? Efficiency here is defined as both an ease of perception, and an intuitive translation to activity requirements. An encompassing question is what can be gleaned about participants’ tacit knowledge of intensity-based sonic feedback from allowing them to provide that feedback themselves? Finally, when it comes to complex, rich, ambiguous sound changes, is there a way of establishing a basis for comparison across participants in the study?

4.2. Exploratory Informant Workshop Design

To create a situated, contextual study for investigating issues of intensity-based gradient sound feedback, we organized an informant workshop centered on a physical game activity. It had to be an informal and engaging activity that at the same time is goal-oriented in an incremental fashion, suitable for a gradient sound feedback, and one whose rules are formalized enough to allow systematic observation and analysis of interactions. To meet these objectives, we devised a geography guessing game with secret locations hidden on a physical poster-size map of the world. In this turn-taking game, one player decides upon a secret location and helps the other players find it. A tangible pointer device was moved across the map to seek the hidden goal. In order to make sound the only form of feedback, we introduced three different instruments used to sonify that game task’s progress (see Table 1).

Participants were explained the “hot and cold” game analogy and told to use only sound generated by one of the instruments of their choosing to help their team mate find the secret location. The instruments were musically simple to mitigate skill and experience factors, and each had a distinctly different sound quality and intensity constraints (e.g. while a kazoo could perform obvious variations in pitch, tempo and timbre, claves and shakers have virtually no variation of pitch, only of amplitude and tempo). We saw this low-tech, participatory format as crucial to our core research question in this study, which was – how would people both generate and interpret sound intensity gradients in relation to a goal-oriented embodied activity?

4.3. Advantages and Disadvantages of this Approach

There are disadvantages to our qualitative approach to auditory display design. For example, a low-tech workshop by definition means a lack of true consistency in the feedback – every time it is provided, it is slightly different, because of variation within and across individuals’ responses. In addition, because engagement is an issue in the context of an informal activity (game or non-game), there is the added complexity of designing a suitable workshop activity. There is rarely time for repetitive trials of the same type of feedback being provided to each participant. Therefore, in our case, where we have three options for providing feedback in addition to several ways of representing intensity, most quantitative measurements that are taken, such as time of completion, could not reliably be used to reflect valid difference between sound feedback approaches or individual perception.

Audio Display	Approach to Intensity	Polarity
Kazoo	Pitch Shift (Complex Tone) + Amplitude	Positive
Clave sticks	Tempo Shift + Amplitude Clean Timbre	Positive
Egg Shakers	Tempo Shift + Amplitude Rough Timbre	Positive
Triangle	Confirmatory Feedback	N/A

Table 1: This table shows the three instruments used in the workshop and their respective intended approaches to intensity

On the other hand, this participatory informal approach can provide surprising depths of qualitative information, revealing important and at times salient patterns of user interactions in relation to directive sonic feedback. Team members represent intensity to each other, creating externalizations of tacit, latent ideas about sound intensity across all the relevant elements of our framework – scaling, polarity, approach to intensity gradient, spatialization and embodied interaction.

4.4. Research Framing and Structure

In order to make explicit our own concepts and expectations of sound intensity feedback, we created a framing schema for analysis of the workshop data. The schema was partially based upon a similar study that was conducted with children. From it, we took the interaction patterns that proved most important and organized them in order to see if they surface again in this iteration of the workshop. First we identified the *sound types* (instruments) and sound *parameters* (pitch, amplitude, tempo) of interest with regard to the musical instruments provided (Table 1). Of course, users were free to deviate from our intended approaches and descriptions, but we thought that having a pre-existing schema would facilitate identifying these instances and analyzing them later. In addition, we externalized the mappings between the 2-D geography map artifact and sound intensity feedback. Specifically, physical *distance* from the pointer device to the secret location at any one time of the game is naturally mapped to the gradient sound intensity. If one is far away from the goal, intensity is low, if they are close to the goal intensity is higher.

At the same time, the concept of *trajectories* (pointer leading towards or away from the goal) was identified as another useful concept, as it functions differently than distance. As exemplified in Fig. 1, Trajectory 1 starts far from the goal, yet moves towards it, so feedback should intensify, while Trajectory 3 starts closer, yet moves away, so feedback should de-intensify, even though it has started from a higher position on the gradient.

Participants took turns providing and receiving auditory feedback during the ‘game’ section of the workshop. In addition, they were given a choice as to which instrument to use but no instructions as to how to represent intensity. All instruments were briefly demonstrated to ensure that players knew how to use them. A total of 4 female users participated, ranging in age from 30 to 51. Six iterations of the game activity took place, including two blindfolded turns, initiated voluntarily by two participants. Only one user had previous musical training, and all reported normal hearing.

We used a passive graph plotting exercise in order to explore intensity-based sound that is complex (involving intensity gradients with several types of changes in pitch, as well as amplitude and timbre) and environmental,. Participants were given forms with 16 blank graphs (see Table 3.) and asked to listen to 16 sound excerpts of 10 to 30 seconds each that represent some kind of intensity change.

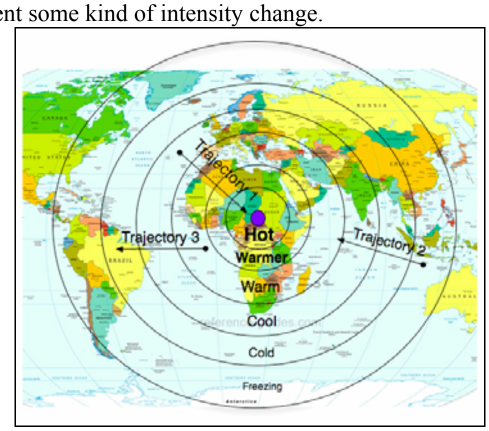

Figure 1. *Model of the world map indicating areas of low, medium and high intensity, as well as three sample trajectories*

They were asked to both graphically plot the change they hear, and write a short one-sentence description of the excerpt. The scaling, polarity and sound content in this section of the workshop were directly borrowed from the auditory feedback model in socio-ec(h)o in order to maintain consistency [24]. Methods of data collection included observation, semi-structured discussion, paper sheets with plotted sound intensity graphs, and video/audio capture.

5. DATA ANALYSIS

One of the biggest challenges with design-based research, especially open-ended participatory workshops and design games, is interpreting the “results” in ways that are meaningful to the design process and/or research question, and reasonably address the issues under investigation. Indeed, in the area of human-centered design there is a certain fascination and enthusiasm for participatory design, because it affords the user the tools and opportunity to express and ideate design [26]. However, it takes careful planning and moderating to make for rich yet useful data.

For these reasons the chosen format of our study was an informant workshop with a more rigid structure and more built-in constraints, permitting fewer distractions. For example, sound feedback was isolated as the only type of guiding response within the game. Also, it was technically impossible for the player to know where the secret location was without feedback, since the other participant decided it on the spot. However, challenges of interpretation still remain in such workshops, and the data analysis often requires creative approaches to representing and deriving useful patterns from the rich qualitative data. Video analysis, audio analysis, conversation analysis, in addition to observational notes and semi-structured discussions make up some of the common methods in the area of *interaction analysis* [29].

In a design context, the goal of conducting such workshops is to further a design outcome such as an artifact or system. In the context of our design research, the system and the workshop goals are to uncover tacit perceptual experiences related to intensity-based sound feedback, in order to eventually inform the design of embodied learning environments. In the current case study, the main focus was on analyzing interactions through video annotation (looking at pointer movements in response to the sonic feedback), interpreting the informal discussions log, organizing and comparing sonic graphs, and analyzing and comparing intensity-based audio feedback (the actual sound tracks from user-generated sonic responses).

5.1. Workshop Game Analysis

In the geography guessing game, several categories of interest were identified for analysis – the user-generated data-to-parameter mapping, scaling, polarity, and chosen sound type. The distinction between location/distance versus trajectory-based feedback is also an important one, because it illuminates whether users understand the concept of a progress gradient. The idea is that there always IS a gradient present, and it just intensifies or de-intensifies according to the physical location of the cursor in relation to the desired end. The gradient should also start at a sound intensity reflective of the initial position of each trajectory. The ideas of location and trajectory did indeed describe a pattern that was observed in the user performances in this iteration of the workshop. Specifically, they tended to provide relatively unchanging, medium-level feedback while the other player was far from the goal, and rapidly intensifying feedback within a small vicinity of the goal (see Table 2). In our interpretation from both the previous workshop and the current one, users tended to focus more on the physical location, rather than on the more abstract idea of adhering to a progress gradient – how close the player is from the goal, versus, where the player is on the map.

There appeared to be an interesting combination of location and trajectory-based approaches to feedback. This was different from the first workshop in that there were more distinct gradient levels of feedback, ranging from complete silence, to low-mid, high-mid, high and extremely high intensity (see Table 2). Also, all players insisted on starting in the middle of the map and performing the intensity gradient from there. It is interesting to note that even though silence was never discussed by the researchers as a form of feedback in either workshop, both sets of participants from our previous and the present workshops intuitively understood it as a default. One user even remarked that “it is really clear when you are not close, the instrument doesn’t make any sound” and so she suggested making a constraint that “silence isn’t a sound that you can do” so that it would push people to provide more specific feedback. The same user remarked during the blindfolded session that she really had to rely solely on the sound because she didn’t even have the visual orientation of the map. This, she described as a much more rewarding and sensory-rich experience. Yet, curiously, she did maintain the right trajectory from the start and quickly found the goal.

In the second blindfolded session one player did not start off right and was constantly sidetracked toward the wrong direction. It became evident just how much the provider of sound feedback struggled to support her quest towards the goal. For our purposes, this was arguably the most interesting iteration of the game, not only from the perspective of user performance with intensity-based sound feedback but also from the perspective of tacit user-driven approaches to intensity-based sound. In order to analyze this rich episode of data, we used a video annotation approach focused on both the trajectory movement in time (video), as well as on the intensity-based sonic feedback in time (audio). In addition, we created a modified scenario-based account of this interaction [27] to generate a story retelling of both sides of the exchange so that relevant issues might surface. Audio files were created from all six instances of the game, with specific attention given to the graphical waveform patterns generated in order to hear and see how scaling was used by different players and whether it went roughly through the same degree of intensity or not.

Similar to the past workshop with children, scaling varied to some degree but overall seemed to reach a maximum of a certain capacity, closely related to human abilities – one could only shake a shaker so fast and loud. Again, whether intuitively, or also because of physical effort (the faster you hit the sticks, the louder it gets), the rate of change in both pitch and tempo, was always tied to amplitude as well. This finding, together with the strong affinity towards silence [no feedback] as a default state aligned with research in psychoacoustics, showing that amplitude is the strongest cue to sound change/intensity [21]. User-generated polarity, as with workshop one, was also always positive – again demonstrated by research in data sonification and contemporary auditory perception.

#	Audio Display Device + Approach to Intensity	Approach to Intensity -Scaling (Feedback Progression)
1	Shaker (Tempo and Amplitude)	No Sound – Faster, medium Sound – Very fast, loud sound
2	Clave sticks (Tempo and Amplitude)	No Sound – Medium tempo – Faster, louder – Very Fast/loud
3	Kazoo (Amplitude reported) – pitch and tempo observed	No Sound – bursts of tonal sounds (high and low pitch) – Loud, droned out high pitched sound
4	Clave Sticks (Tempo observed)	Slow – Medium – Faster – Very Fast
5	Shaker (Tempo)	No sound – Slow – Fast – Very Fast
6	Clave Sticks (Tempo and Amplitude)	No Sound – Very Slow – Medium/louder – Very Fast/loud

Table 2: Coded results outlining participants’ observed and/or reported use of the audio feedback devices and approach to representing intensity

5.2. Sonic Graph Plotting Analysis

The passive graphic plotting task was an effective way to isolate and test out a number of key sound characteristics, intensity gradients, sound types and ranges in a way that still allowed users to externalize and represent their perception and understandings in an intuitive way. 16 sound excerpts were tested. Sounds started from simple/sine tones with straightforward parameter changes and moved towards complex, everyday sounds with more subtle and compound parameter changes. The first two excerpts were sine tones, and they were partially used as a way to establish 'sound intensity competence' – which all users seemed to possess (see Table 3) as all correctly identified a rising and falling sine tone. The rest of the sounds were designed to be more ambiguous, involving degrees of greater or less change, and sound types that varied from musical, abstract to environmental. More than anything, this graph plotting sound identification portion of the workshop was conducted as proof of concept for exploring perception of complex changing sound. We saw graphing, rather than choosing a range or number, as a more direct way of getting at users' tacit perceptual experience. The rigid structure of a graph makes it possible for different users' graphs to be compared directly. While we did not know what kind of results we might get from this, the examples presented in Table 4 make clear that even with a small sample and noticeable variation in drawing styles, definite similarities and common patterns can be discerned when comparing user representations of changing intensity-based sound.

. However, it was still important that users had some space to write a short descriptor for each sound, as some sounds might have been too complex to graph, and verbal description could better represent what participants heard. This turned out to be a wise decision, because when reading and comparing the graph-plots and transcripts, it was discovered that users often drew different-looking graphs, but expressed similar experiences through words, and the other way around.

Because of the seeming gap between psychology studies of auditory perception and their application to context-sensitive perceptual frameworks for sound, making sense of our data is still at its early stages. Following the methods of interaction analysis we have constructed a visual comparison schema that includes the actual (intended) intensity of the sounds, and a composite of each participant's plotted graph of it (Table 3). We provisionally suggest that this test is in itself a valuable contribution towards developing non-traditional forms of data representation for exploring and analyzing changing everyday sound that often reveals aspects and perspectives of research that could otherwise be missed.

Type and Degree of Change (intended)	**Participant- Generated Graphs (interpreted)**
Sine tone at 100Hz rising steadily to 440Hz 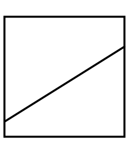	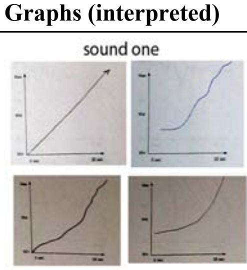
Pulsating sine tone of 220Hz increasing in tempo from 12% to 87% 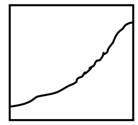	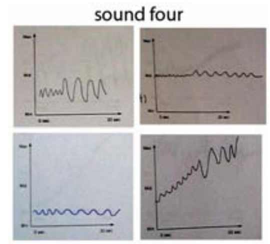
Fire sound through a low-pass filter. Starts muffled and quickly is unfiltered (0% to 100%) 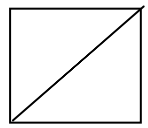	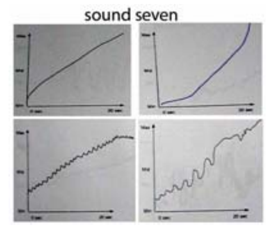
Ticking clock sound through a low-pass filter – starts crisp, then becomes muffled then crisp again. 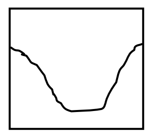	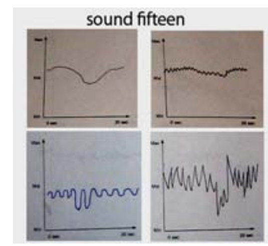

Table 3. This table shows a few examples from the sound plotting section of the workshop. On the models to the left, the X-axis represents time in seconds, while the Y-axis represents intensity from 0% to 100%.

6. DISCUSSION

Using small-scale participatory workshops may be the long road to constructing a functional model of situated human perception of complex, changing sound in contexts of engagement and active, embodied learning. Yet such small-scale situated explorations may offer optimal conditions for achieving ecological validity in this area of auditory display design. Thus, we hereby focus our efforts on suggesting ideas and ways of approaching design and research problems with sound feedback for embodied learning, as well as innovative ways for analyzing the rich, situated data generated as a result.

Rather than attempting to advance general claims about the salient efficiency hierarchy of timbre-based versus pitch or rhythm-based feedback, we have drawn lessons from our previous workshop-based trials and tried in an increasingly nuanced and fine-grained way, to build upon these. What has proved most interesting already has been the process of allowing users to generate intensity-based sonic feedback themselves, in the context of a playful activity. That is where novel insights were suggested with respect to the latent perceptual-conceptual mappings between activity, representations and intensity-based sonic feedback. The research challenge is finding analytical approaches that enabled the parameters of interest to be teased out from the record of that activity. As mentioned earlier, analysis of design-based research workshops for interaction design is largely understudied and lacks stable guidelines, even though, we suggest, it affords significant opportunities for novel results. We hope this study contributes to the design of innovative ways of facilitating visual analysis of user-generated data (in our case audio and video footage). We propose that seeing things in alternative ways can sometimes reveal or elucidate issues previously overlooked, or that simply could not be discovered using other, more conventional methods of analysis. We are especially interested in analyzing user-generated audio captures in more detailed ways, breaking down their temporal and frequency structures and comparing these patterns across cases. The process of examining sound through user-generated graphs is also promising. As well, we see considerable potential in developing a modified scenario-based approach with which to explore retelling of research activities as narratives in order to capture descriptive, yet also intuitive, analytical and observational features of the interactions. Involving users in such ways falls within and furthers design-based research methods.

7. CONCLUSION

The conceptual and operational model here described, building as it does upon prior participant workshop-based studies for researching situated auditory display design offers promising directions for future explorations of information-rich, ambient sound feedback for multimodal, embodied learning environments, as well as suggesting empirically well-grounded bases for refining questions and developing more focused and nuanced approaches to documentation and analysis.

8. REFERENCES

[1] Sanchez, J., Flores, H. AudioMath: blind children learning mathematics through sound. In Proc. ICDVRAT 04, ACM Press (2004): 183-89.

[2] Parente, P. and G.Bishop. BATS: The Blind Audio Tactile Mapping System. In Proc. ACM SRC ACM Press (2003).

[3] McElligott, J. Designing Sound Tools and Toys for Blind and Visually Impaired Children. In Proc. IDC 2004, ACM Press (2004), 65-72.

[4] Takahata, M. et.al. Sound Feedback for Powerful Karate training. In Proc. NIME 2004, (2004).

[5] Wakkary, R., Hatala, M. Lovell, R, & Droumeva, M. (2005). Ambient intelligence platform for physical play. In *Proceedings of the 13th annual ACM international conference on multimedia* (pp.764-773). New York: ACM Press.

[6] Nesbitt, K. V. and S. Barrass (2002), "Evaluation of a multimodal sonification and visualization of depth of market stock data.," presented at International Conference on Auditory Display, Kyoto, Japan, 2002.

[7] Butz, A., Jung, R. Seamless User Notification in Ambient Soundscapes. In Proc. IUI 05, ACM Press (2005) p.320–322.

[8] Dourish, P. (2001) Where the Action Is: The Foundations of Embodied Interaction. Cambridge: MIT Press.

[9] Sharlin, E. et al. (2004) The Tangible Pathfinder: Design of a Wayfinding Trainer for the Visually Impaired. In Proc. Graphics Interface, 2004.

[10] Truax, B. (2001) Acoustic Communication. Ablex Publishing. 2nd Ed.

[11] Schafer, R.M. (1977) The Tuning of the World. New York: Knopf Press.

[12] Walker, B, and G. Kramer, Mappings and metaphors in auditory displays: an experimental assessment," In Proc. ICAD'96, (1996).

[13] Adcock, M., Barrass, S. Cultivating Design Patterns for Auditory Displays. In Proc. ICAD 04 (2004).

[14] Shinn-Cunningham, Barbara and Antje Ihlefeld. (2004). Selective And Divided Attention: Extracting Information From Simultaneous Sound Sources. In Proc. ICAD 04, Sydney, Australia (2004).

[15] Hunt, Andy, & Hermann, Thomas. (2004). The Importance of Interaction in Sonification. In Proc. ICAD 04, Sydney, Australia (2004).

[16] Shinn-Cunningham, B., Ihlehem, A. Selective And Divided Attention: Extracting Information From Simultaneous Sound Sources. In Proc. ICAD 04 (2004).

[17] Droumeva, M., & Wakkary, R. (2006). Sound intensity gradients in an ambient intelligence audio display. In *CHI 2006: Conference on human factors in computing systems* (pp.724-729). New York: ACM Press.

[18] Gaver, W., R. B. Smith, and T. O'Shea, Effective Sounds in Complex Systems: The ARKola Simulation, In Proc. CHI'91 ACM Press (1991).

[19] Cook, Perry. (1999) Music, Cognition and Computerized Sound. MIT Press

[20] Neuhoff, J. Ed. Ecological Psychoacoustics. Boston: Elsevier Academic Press. (2004).

[21] Ballas, J. Common Factors in the Identification of an Assortment of Brief Everyday Sounds. Journal of Experimental Psychology: Human Perception and Performance, 19(2), (1993) 250-267.

[22] Droumeva, M., Wakkary, R., Participatory Design for an Ambient Intelligence Audio Display, *In Proc. International Conference on Audio Display 2006*, pp. 36-43.

[23] Wakkary, R., Hatala, M. Lovell, R, Droumeva, M., Antle, A., Evernden, D. & Bizzocchi, J. (2005). socio-echo: Ambient intelligence and gameplay. In *DIGRA 2005 International conference: Changing views: Worlds in play*.

[24] Carroll, J.M., Chin, G., Rosson, M.B. and Neale, D.C., The development of cooperation: Five years of participatory design in the virtual school. In Proc. in DIS 2000, ACM Press, 239-251 (2000).

[25] Sanders, E.B.-N. Collective creativity. Loop: AIGA Journal of Interaction Design Education, 3. 2001.

[26] Svanæs, D. and Seland, G., Putting the users center stage: Role playing and low-fi prototyping enable end users to design mobile systems. *In Proc. in CHI 2004*, (2004), ACM Press, 479-456.

[27] Barab, S. and Squire, K. (2004) Design-based research: Putting a Stake. *The Journal of the Learning Sciences* 13(1): pp.1-14.

[28] Henderson, A. Interaction Analysis: Foundations and Practice, The *Journal of Learning Sciences*, 4(1): pp. 39-103 (1995).