

PARTICIPATORY DESIGN RESEARCH METHODOLOGIES: A CASE STUDY IN DANCER SONIFICATION

Steven Landry, Myounghoon Jeon

Mind Music Machine Lab
Michigan Technological University
Houghton, Michigan, 49931
{sglandry, mjeon}@mtu.edu

ABSTRACT

Given that embodied interaction is widespread in Human-Computer Interaction, interests on the importance of body movements and emotions are gradually increasing. The present paper describes our process of designing and testing a dancer sonification system using a participatory design research methodology. The end goal of the dancer sonification project is to have dancers generate aesthetically pleasing music in real-time based on their dance gestures, instead of dancing to pre-recorded music. The generated music should reflect both the kinetic activities and affective contents of the dancer's movement. To accomplish these goals, expert dancers and musicians were recruited as domain experts in affective gesture and auditory communication. Much of the dancer sonification literature focuses exclusively on describing the final performance piece or the techniques used to process motion data into auditory control parameters. This paper focuses on the methods we used to identify, select, and test the most appropriate motion to sound mappings for a dancer sonification system.

1. INTRODUCTION

Evidence supports that multimodal interactions increase user engagement with novel interfaces [1]. Therefore, sonification can buttress the connection between the receiver and the information, exploring a new form of art by a synesthetic combination of music and dance. Interactive sonification can be defined as “the use of sound within a tightly closed human-computer interface where the auditory signal provides information about data under analysis, or about the interaction itself, which is useful for refining the activity” [2]. As an interactive sonification technique, parameter mapping [e.g., 3] has often been used, where data features are arbitrarily mapped onto acoustic attributes such as pitch, tempo, timbre, etc. From this background, we have devised a novel system, immersive Interactive Sonification Platform (“iISoP”) for location, movement, and gesture-based interactive sonification research, by leveraging the existing Immersive Visualization Studio (IVS) at Michigan Technological University. The iISoP has been developed for multi-disciplinary research in a verity of fields such as data sonification, gesture interfaces, affective computing, and digital artistic performance. The present paper discusses issues, considerations, and strategies currently implemented in the iISoP's dance-based sonification project, in hopes to spur

discussion of applications of more artistic interactions in the sonification community. The selection and fine tuning of motion-to-sound parameter mappings are at the core of any sonification project. Choosing and evaluating these mappings require a network of interdisciplinary team members, each with specific goals and design philosophies that may not always align together. How these decisions are resolved and evaluated are highlighted through a case study in dancer sonification.

1.1. Dancer Sonification

Under normal dance circumstances, the choreographer designs the dance to match specific music. To refer to this type of connection between visual and audio content in multimedia, the term “synchresis” was recently coined [4]. Certain gestures and emotions are utilized to match with specific movements of the musical piece. In the dance-based sonification project of the iISoP, the reverse process is implemented. Music is generated in real-time based on the dance to increase the amount of synchresis between the visual and auditory characteristics of the entire dance experience. The end goal of the dance-based sonification project is to have dancers generate aesthetically pleasing music in real-time based on their dance gestures, instead of dancing to pre-recorded music. The generated music should reflect both the kinetic activities and affective contents of the dancer's movement. The dancer begins to dance, and the sonification system interprets the movements and generates music. The generated music, in turn, influences the way the dancer dances, which is again sonified, leading to a closed loop between the dancer and the system in an interactive manner. To this end, we have collaborated with multidisciplinary teams, involving cognitive scientists, computer scientists, sound designers, dancers, and performing artists.

This dancer sonification project falls in line with previous projects on dance sonification such as the DIEM digital dance system [5], The MEGA project [6], and David Rokeby's Very Nervous System [7]. The iISoP's approach to dance sonification differs from these past projects in several ways. Our goal for the iISoP is to generate aesthetically pleasing music that is composed of multiple layers or streams of instrumentation. Multiple streams are important to build the body of a musical piece, an important aspect for the immersion of both the dancer and audience. An additional task that previous versions of dance sonifications ignored is affect detection of the dancer, and synthesis of affective content of the gesture sonification that reflects the current state of the dancer.

As with any design research project, the critical aspects to be documented and reported are the methods for which the design is constructed.



This work is licensed under Creative Commons Attribution – Non Commercial 4.0 International License. The full terms of the License are available at <http://creativecommons.org/licenses/by-nc/4.0/>

<https://doi.org/10.21785/icad2017.069>

2. SYSTEM CONFIGURATION

The immersive interactive sonification platform (iSoP) is an interactive multimodal system. Figure 1 shows a conceptual diagram of the iSoP system configuration.




Figure 1: Architecture and data flow of the iSoP system.

The iSoP features a Vicon tracking system utilizing 12 infrared cameras that track specific reflective objects that are strapped to the user's limbs (e.g., wrists and ankles) via Velcro bands. The visual display wall consists of 24 42" monitors controlled by 8 computers that display representations of the tracked objects in real-time. Position, velocity, acceleration, time, proximity to other objects, and holistic affective gestures are recorded and analyzed to generate appropriate sounds (speech, non-speech, music, etc.) based on our own sonification algorithms programmed in Pure Data (real-time graphical dataflow programming environment) [8]. Motion data can also be routed through Wekinator (an open source software tool for real-time interactive machine learning) [9] for machine learning recognition of body postures and gestures. MIDI (Musical Instrument Digital Interface) and OSC (Open Sound Control) messages can also be sent from Pure Data to Ableton Live (music software for MIDI sequencing and music production, creation, and performance) for additional sound synthesis.

3. METHODS

Below are the methodologies in chronological order that we have employed in an attempt at creating an interactive and musically expressive dance-based sonification system.

3.1. Collaboration with performing artist Tony Orrico

As a testbed of our visualization and sonification system, we invited an artist to perform in the iSoP. Tony Orrico is a Chicago based performing artist known for creating large geometric pieces (e.g., Penwald Drawings), using his entire body as an instrument in artistic expression [10]. Orrico demonstrated one of his penwald drawing pieces while wearing sensors that made real-time visualization and sonification. Tony, being a mainly visual artist, had little to contribute to the "sonification" design of the performance, which gave the sound designer full autonomy to choose and implement all parameter mappings. The goal for the sonifications was similar to the goals of the visual presentation: to add a technological aesthetic to the performance piece by reinterpreting and representing the analog movements digitally in real-time.

The sound designer programmed four arbitrary melodies (MIDI format) approximately 1-2 measures long. These MIDI melodies were sent to a "digital bell" sounding MIDI instrument. The instrument and melodies were chosen by the sound designer to convey a particular aesthetic, one that

invokes imagery of meditative prayer bells in a monastery mixed with an electronic synthesizer. Melodies were played in an arbitrary order, and the speed of playback was determined by referencing the current velocity of the artist. The purpose of this mapping was to convey to the audience the relationship between the artist's physical movement and the sonic feedback.

For one performance, the artist Tony laid face down on a piece of paper holding graphite pencils in both hands. He pushed off a wall, jetting himself forward on top of the piece. He dragged his graphite pencils along with him; as he writhed his way back to the starting position over and over again, he left behind himself a pictorial history of his motion. While he was drawing these pieces on the paper canvas, his movements created another drawing on the virtual canvas (display wall of monitors) as well as the previously described real-time sonification.

3.1.1 Evaluation

Informal feedback was collected from audience members after the performance was complete. Unfortunately, the audience was generally not impressed with the sonification aspect of the performance. Audience members felt the sonification added little if anything to the overall experience, and it failed as an auditory representation of the performer's physical movements.

On reflection, the sonification "failed" on an implementational level. Melody playback rate was only updated the instant the previous melody finished playing. For instance, if the system happened to update melody speed during a portion of the routine with low velocity, the next melody would be played at a very slow rate (1/4 note = whole notes), which could last for almost 30 seconds. This slow melody would spend 30 seconds describing one instant of the performer's past velocity, growing more irrelevant to the current state of the performer as time passes. If the system happened to update melody speed during the short high-velocity portions of the performer's routine, a quick short melody was triggered, which in the designer's opinion successfully described the movement sonically. Unfortunately, this occurred very rarely. Overall, it was this lack of synchronicity between the activity of the performer and sonic feedback that led to the poor reviews of the sonification. The most obvious movements to the audience (Orrico pushing/jetting from the wall) were often completely ignored by the sonification. This suggested that future versions of the dancer sonification system should include more continuous mapping (rather than triggering discrete melodies) to ensure more synchronicity between motion and sound.

3.2. Dancer interviews

We knew that in order to develop an interactive dance-based sonification system, our background in usability, interactivity, and audio design could only take us so far. With no one on the project having any formal dance training, it was critical to incorporate feedback from domain experts and end users. To this end, we conducted a number of interviews with expert dancers to 1) gather system requirements, 2) evaluate the current versions of our system, and 3) generate novel and intuitive interaction styles and sonification techniques.

Six expert dancers were recruited through local dance performance schools and the local university's Visual and Performing Arts Department. All dancers had at least 10 years of professional dance training. Each semi-structured interview was performed individually, lasting from one to two hours.

The first section of the interview revolved around the expert dancer describing what they would imagine a dancer sonification system to be. This was done before the dancer experienced the current sonification scenario to avoid any anchoring bias. The next section involved the dancer interacting with the system for around 15 minutes while describing their impressions in a “think aloud” fashion. The final section of the interview included a brainstorming session for suggesting modifications and additions to the system, as well as potential applications for the system in other domains.

One interesting theme that came up multiple times through the expert interviews was the importance of valuing the visual aesthetic of the dance over the aesthetic of the sonifications. This has implications over how much control the dancer wishes to have over the sonifications. For instance, dancers would not want to contort their body into odd shapes just to achieve a desired sound. Dancers should also not have to consciously consider every aspect of the sonifications when determining which gesture or posture to perform in sequence. One expert dancer explicitly said “I want 50% of control over the music so I can concentrate on the dance as much as possible”. This would require a certain amount of automation on the system side to produce novel and interesting music describing the motion and emotion of the dancer, which accords to the previous experiment [12]. This was in direct conflict with the sound designers associated with the project, who imagined having complete control over every aspect of the sound generation. Musicians place little to no value on the visuals of the gestures, placing all value on the acoustic properties of the sound. From an HCI research perspective, the value is placed on how the user’s performance and impression change when the evolutionarily established feedback loop between the dancer and sound is augmented or reversed using technology. In general, each stakeholder has individual goals and philosophies for the project that are at best loosely related, and at worst completely contradictory.

3.3. Visual and Auditory Stimuli Collection

After conducting six interviews, we aggregated general concepts for what expert dancers envisioned how the system should behave. It should first interpret the gestures and affective content of the dance, then create music describing that information. In order to “teach” our system how to perform these tasks, we first had to investigate how humans would accomplish this task. We needed to identify heuristics human composers use to interpret and sonify the motion and emotion of a dance performance. To identify these heuristics, we conducted a small study to collect and analyze visual and auditory stimuli. This study had two goals: 1) to see how and how well non-experts could detect emotion from dance gestures, and 2) to see what type of music or sounds human composers would use to describe dance gestures.

To address the first goal, we invited two expert dancers to submit video recordings of themselves dancing to popular music. These two dancers also participated in the initial interviews. The dancers picked popular songs that represented a particular basic emotion (Anger, Happiness, Sadness, or Content), and performed a dance routine that attempted to portray that particular emotion visually. We then recruited 10 novice participants to watch muted versions of the videos and classify each with the emotion the dancer was attempting to portray. To address the second goal of this exercise, we recruited a music composition class consisting of 10 amateur composers to sonify muted versions of the dancer videos. We

gave three specific instructions to the composers. Composers where to: A) re-imagine and recreate the music that the dancers were originally dancing to, B) score the video as if for a film, focusing on capturing the overall mood of the dancer, and C) compose a collection of sounds that describe the kinetic movements of the dancer.

The results of the “guess the emotion” portion of the study suggested that it can be difficult for people to express and interpret emotion through dance gestures alone. There was very little agreement amongst the responses, and self-reported confidence scores were very low. This could be due to a number of factors, but the two most likely explanations of the low accuracy and agreement are 1) communicating emotion through dance is difficult, or 2) non-dancers have difficulty interpreting the intended emotion from dance gestures. Overcoming these obstacles will be critical for embedding automated affect detection algorithms in the iSoP system. The results of the audio stimuli collection portion showed just how infinite the problem space is when considering what type of motion to sound parameter mappings could (or should) be implemented in our dancer sonification system.

Some parameter mapping sonification strategies were consistently used in the majority of audio submissions. Dance gestures that involved rising limbs (raising an arm or leg) were often accompanied with melodies that increased in pitch, and vice versa. Larger body movements were often paired with “larger” sounds (e.g., polyphonic chords, multiple instruments, increase in volume, etc.). Speed of dance gestures was also commonly paired with the speed of the melody (subdivision rate, not BPM of the song). As a note, the project’s sound designer was solely responsible for identifying motion-to-sound parameter mappings used in the compositions. This introduces a bias in the type of mappings extracted from the submissions. For instance, mapping height to pitch and velocity to speed was already the intention of the sound designer all along. The same biases certainly unintentionally might filter the information extracted from the expert interviews as well, as the designer could not fully compartmentalize their own goals and philosophy from the interviewee.

3.4. Three new dancer sonification scenarios

We wanted to design a few sonification scenarios leveraging these general strategies used by the human composers from the stimuli validation study. In order to move towards more continuous parameter mapping, we incorporated the real-time graphical programming environment Pure Data into the iSoP architecture. Pure Data afforded us the ability to program virtually any algorithm for real-time parameter mapping sonification. However, designing aesthetically pleasing instruments in Pure Data is time consuming for even the most proficient programmer. In order to leverage the expressivity and control of sound that more conventional DAWs (digital audio work stations) afford to the non-programming population, we included Ableton Live as an alternative means to design and play more aesthetically pleasing instrument sounds.

Two common algorithms we programmed in Pure Data attempted to translate (or map) height to pitch, and velocity to subdivision of generated melodies. For those interested, pictures of the Pure Data subpatches implementing these algorithms are presented below.

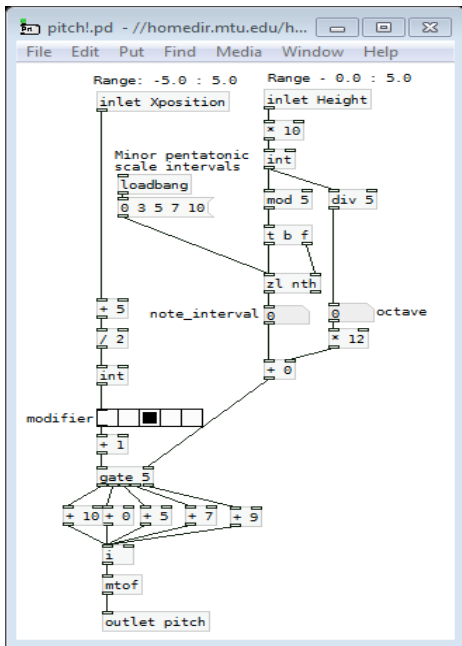


Figure 2: Sonification algorithm for translating position of a tracked object into a MIDI pitch.




Figure 3: Sonification algorithm for translating velocity of an object to when and how long to play a note within a given measure.

The first of the three newly created scenarios (“A”) focused around a theme of using a user’s body as an instrument. Each hand controls independent instruments (melody and percussion). There is a direct mapping between movement speed of that hand and the volume/rate of the arpeggiator for that hand’s instrument. Note pitches for the tones are rounded to the nearest note in key, and the onset/duration of notes are quantized in time to the nearest 32nd note subdivision of the tempo. Similar time quantization is used for the percussion instrument using a Euclidean rhythm generator, where the tracked object’s current speed determines how many

percussion hits are equidistantly distributed across a one measure phrase. The percussion instrument consists of synthetic hi-hat clicks and a bass drum sample. Hand velocity control for the bass drum is scaled down to 1/3 of the rate of the hi-hat clicks to create a syncopated drum rhythm. To provide constant timing cues, a synthetic snare drum is constantly playing on beats two and four of the measure independent of the user’s movements. All variable scaling and sound production are done through Pure Data .

The second scenario (“B”) focused around a theme of using a user’s body as a DJ’s MIDI controller. A very simple 4 measure musical loop was created as a set in the Ableton Live. A number of motion variables were scaled to MIDI range (1-128) using a custom Pure Data patch and routed to through Ableton’s MIDI mapping functionality. The user can control a number of parameters controlling the playback of certain instrument tracks or an audio effect applied to the master output. For instance, the right hand’s height controls the amount of filter added to a distorted bassline, and the distance between the two hands determines the cutoff frequency of a low pass filter applied to the entire loop playback.

The third scenario (“C”) was a kind of hybrid of the first two themes, where different aspects of the body’s overall shape is mapped to a 3 dimensional fader slider controlling the volume balance between 8 pre-made musical loops. Eight musical loops were collected from an online database (all 120 BPM, in the key of C minor, with a length of one, two, or four measures). The musical loops were loaded into a 3D fader object in a custom Pure Data patch for synchronized playback, where each corner of the cube corresponds to one of the eight musical loops. The distance of current position of the fader slider to each of the 8 corners of the cube determines the volume of each of the corresponding musical loops. Six different body shapes (described by distances between the tracked objects) were mapped to the min and max of each of the 3D slider’s position variables (X, Y, & Z) using Wekinator. As the user dances or changes poses, the three dimensional fader raises or lowers the volume of each of the 8 musical loops, creating interesting combinations of melodies and rhythms. Note that a sound designer oversaw and configured sonifications of all three scenarios and so, overall sound quality could be similar across the three scenarios.

3.5. Dancer sonification scenario evaluation

In order to evaluate and compare these three newly created scenarios, we conducted a study to evaluate the overall system performance and sonification strategies. Specifically, we wanted to investigate what effect the different interaction styles for each scenario have on user impressions of flow, presence, and immersion in the virtual environment.

Twenty-three novice dancers ($M_{age} = 20.3$, $SD_{age} = 2.1$) participated in the study. All participants were recruited from the local university’s undergraduate psychology program in exchange for course credit. Eleven participants reported some musical training, and six participants reported some (below 4 years) of formal dance training. Each participant experienced each of the three sonification scenarios for roughly five minutes. This involved the participant exploring and interacting with the system through improvisational dance. Following each scenario, the participant filled out a battery of questionnaires including measures of flow, expressivity, and immesiveness. Participants were also instructed to try and discover and report what motion-to-sound mappings were present in that particular scenario.

Scenario A was reported to have the most “discoverable” or “intuitive” motion-to-sound mappings. Most participants were able to discover at least three of the motion-to-sound mappings regardless of their dance or music demographic backgrounds. Reviews for the overall aesthetics of the sonifications were mixed. Many participants reported the ability to control aspects of the sound that they in reality could not.

Scenario B consistently scored the lowest on the majority of the scales. Many participants reported that the interaction style was confining, not intuitive, and did not encourage the exploration of novel movements. Musicians (especially those who had some experience with digital audio workstations) were more likely to enjoy scenario B and discovered more mappings than non-musicians.

Scenario C was by far the most preferred scenario of the three, and participants suggested it had the most potential for artistic performance applications of the three. Scenario C was also believed to have the most amount of features, even though technically it had the least amount of motion-to-sound mappings. A few participants reported that the interaction style in C was “gratifying”. Most participants also mentioned that scenario C’s sonifications were the most pleasant sounding of all three scenarios. Participants reported that scenario C’s sonifications worked “as a sound representation of the user’s movement”, the best out of the three scenarios. This was counterintuitive to the designer’s expectations, as scenario A was designed to have the most obvious 1:1 mappings between movement activity/location to sound. Scenario C also scored highest with respect to the “the sound helped me understand my movements better” agreement statement.

An interesting finding is that participants often perceived more control of the music than they actually had. For instance, a participant with 4 years of formal dance training reported that he thought he could trigger the synthetic snare drum in scenario A with a sharp deceleration of body movements. In actuality, the snare drum constantly played on beats two and four regardless of user behavior. This was a feature designed to provide familiar temporal cues to the dancer with respect to the tempo and beat of the measure. However, since dancers have been trained to synchronize their movements to these temporal cues, the participant naturally (or unconsciously) synchronized his movements to the automated snare drum. He mistakenly attributed this temporal “coincidence” between motion and sound as a causal relationship. This observation raises additional research questions, such as “what other learned dance behaviors can we leverage to facilitate a richer interaction between user and system?”.

Although scenario B was made by a musician for a musician, participants with musical training still preferred the other two scenarios as a whole. Perhaps, a few of the mappings in scenario B were too subtle for non-musicians to notice. In the future, more obvious movements should correspond to more obvious changes in the sonic feedback. Control metaphors used by the designer to control the sound had to be explained to the participants, which suggests these metaphors are not generalizable to others. For instance, the X distance between the hands controlling the low pass filter cutoff frequency was intended to be a metaphor for compressing or stretching the sound as if it was a tangible object.

It was most likely a combination of 1) the clear target goal (isolating an individual loop or achieving a corner position in the 3D fader cube), 2) the challenging method of control through manipulating a body’s overall shape, 3) the continuous audio feedback describing the similarity/distance between the

current and target body shape, and 4) the obvious and rewarding sound produced once the target shape was achieved that led multiple participants to report that scenario C was “gratifying”. Many participants suggested combining aspects of different scenarios for a more expressive performance. Future iterations of the iISO’s dancer sonification phase could combine obvious 1:1 mappings of scenario A and the complex interaction style of scenario C.

In addition to these considerations, more technical aspects of the tracking system need to be revisited. Many of the expert dancers (as well as the non-dancing participants) complained that the objects attached to the ankles and wrists of the user restrict movement, and that more places on the body should be tracked. Before we start adding in more sensors, smaller and more comfortable versions of the sensors need to be designed and tested. The location of hands and feet are only a fraction of the visual information humans use to interpret body posture. Many forms of dance focus on other areas on the body, such as the head, hips, shoulders, elbows, and knees. More data should be collected and used describing the extension angle of joints. There were also struggles with the quality of data from the motion tracking system. Since the dancer’s movements often involve spinning, jumping, rolling, the trackable objects worn by the dancer would often be occluded from the vision of the motion tracking cameras, resulting in a large amount of missing data. We also implemented an instantaneous velocity calculation, which resulted in exaggerated jumps in the reported velocity/acceleration data. We will switch to using a rolling average instead to smooth out the data in future scenarios.

3.6. Dancer Workshops (ongoing)

Another set of scenarios are currently being developed based on the feedback described in the dancer sonification evaluation study. These scenarios will be designed and evaluated during multiple workshops in collaboration with invited expert dancers. Dancers will present during the programming of Pure Data patches to help inform the programmer of appropriate scaling values when translating motion to sound parameters. It is expected that once the dancers have a general feeling of the types of algorithms implemented in Pure Data (through this interactive process), they will be able to suggest more creative potential parameter mappings than in previous brainstorming sessions.

The main direction of the new scenarios is to give the user the ability to control the overall structure and flow of a song, opposed to the static set of instruments featured in the first three scenarios. Since the majority of popular music is structured into repeated sections (intro, verse, chorus, bridge, outro, etc.), giving the user the ability to switch between these sections is another step to accomplishing the end goal of the dancer sonification system. Programmatically, this suggests that sets of premade instruments should be available at all times to the user. The user should also be able to activate, mute, or modify the pre-made instrument sets through specific gestures, locations in the room, or through intervening actions taken by the iISO system based on a rolling average of the “quality of movement” of the dancer. The software “Eyesweb” [13] shows promise in calculating and routing automated “quality of movement” analysis to our sonification software. The quality of motion is based on Laban Movement Analysis, and affords us a much better description of dance gestures than simple velocity and distance calculations.

4. DISCUSSION AND DESIGN CONSIDERATIONS

From all of these works, we have learned how valuable domain expert feedback can be for designing a dancer sonification system. We have also learned how difficult it can be to integrate competing ideas from different stakeholders. We are also starting to unpack exactly how the interaction style and sonification methods can influence users' feelings of presence and immersion in virtual environments. Simply affording the users the ability to control certain aspects of the auditory display does not guarantee interactivity, nor does it guarantee that the users feel "immersed" in the virtual environment. More features and more complex mappings do not equate to "richer" interactivity. Users do not have to completely understand every motion-to-sound mappings in order to express themselves artistically. There are certain aspects of the auditory display that users expect to be able to control, and are disappointed when the system does not conform to their expectations. However, what is perhaps more useful is knowing which aspects of the auditory display can be automated to ensure the music is aesthetically pleasing while still depending on the user's input. These automated strategies alleviate the users' workload to focus on the more creative aspects and dance and composition instead of "trivial" aspects such as specific MIDI pitches and note lengths. We have also learned that the efficacy of different control metaphors is heavily dependent on the user's personal experiences.

Creating a balance between user control and system automation is difficult. Enough automation is necessary to ensure the sonic output of the system is pleasant and structured, like typical popular music. However, embedding too much automation begins to deteriorate the perceived connection between the gestures and music. Giving the user too much control of the sonic output has negative effects on the cognitive flow of the dancer, and the physical flow of the dance performance. A certain amount of stochasticity in the mappings or sonic output may be necessary to keep the music from becoming repetitive. It is important to include what we know about how expert humans compose music (heuristics) in the design of sonification algorithms. Keeping notes in key and using a constant BPM are obvious composition heuristics, as is spreading out audio streams over wide frequency spectrum (e.g., bass, melody, lead, percussion). Designers must keep in mind that the music must still sound musical, and the dance must still resemble dance, otherwise it is no longer a dancer sonification system.

5. REFERENCES

- [1] Hermann, T., & Hunt, A. (2004). The discipline of interactive sonification. In Proceedings of the International Workshop on Interactive Sonification.
- [2] Hermann, T. (2008). Taxonomy and definitions for sonification and auditory display. In Proceedings of the 14th International Conference on Auditory Display (ICAD 2008).
- [3] Dubus, G., & Bresin, R. (2013). A systematic review of mapping strategies for the sonification of physical quantities. *PloS one*, 8(12), e82491.
- [4] Bencina, R., Wilde, D., & Langley, S. (2008, June). *Gesture≈ Sound Experiments: Process and Mappings*. In NIME (pp. 197-202).
- [5] Siegel, W., & Jacobsen, J. (1998). The challenges of interactive dance: An overview and case study. *Computer Music Journal*, 22(4), 29-43.
- [6] Camurri, A., De Poli, G., Friberg, A., Leman, M., & Volpe, G. (2005). The MEGA project: Analysis and synthesis of multisensory expressive gesture in performing art applications. *Journal of New Music Research*, 34(1), 5-21.
- [7] Rokeby, D. (1998). The construction of experience: Interface as content. *Digital Illusion: Entertaining the future with high technology*, 27-48.
- [8] Puckette, M. (1996). Pure Data: another integrated computer music environment. *Proceedings of the second intercollege computer music concerts*, 37-41.
- [9] Fiebrink, R., & Cook, P. R. (2010). The Wekinator: a system for real-time, interactive machine learning in music. In Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)(Utrecht).
- [10] Jeon, M., Landry, S., Ryan, J. D., & Walker, J. W. (2014, November). Technologies expand aesthetic dimensions: Visualization and sonification of embodied Penwald drawings. In *International Conference on Arts and Technology* (pp. 69-76). Springer International Publishing.
- [11] Walker, J., Smith, M. T., & Jeon, M. (2015, August). Interactive Sonification Markup Language (ISML) for Efficient Motion-Sound Mappings. In *International Conference on Human-Computer Interaction* (pp. 385-394). Springer International Publishing.
- [12] Jeon, M., Winton, R. J., Henry, A. G., Oh, S., Bruce, C. M., & Walker, B. N. (2013, July). Designing interactive sonification for live aquarium exhibits. In *International Conference on Human-Computer Interaction* (pp. 332-336). Springer Berlin Heidelberg.
- [13] Camurri, A., Hashimoto, S., Ricchetti, M., Ricci, A., Suzuki, K., Trocca, R., & Volpe, G. (2000). Eyesweb: Toward gesture and affect recognition in interactive dance and music systems. *Computer Music Journal*, 24(1), 57-69.