The Design of a Robotic Marimba Player – Introducing Pitch into Robotic Musicianship

Gil Weinberg
Georgia Tech
840 McMillan St. Atlanta GA 30332
(1)-404-894-8939
gilw@gatech.edu

Scott Driscoll
Georgia Tech
840 McMillan St. Atlanta GA 30332
(1)-404-894-8939
scott.driscoll@gatech.edu

ABSTRACT
The paper presents the theoretical background and the design scheme for a perceptual and improvisational robotic marimba player that interacts with human musicians in a visual and acoustic manner. Informed by an evaluation of a previously developed robotic percussionist, we present the extension of our work to melodic and harmonic realms with the design of a robotic player that listens to, analyzes and improvises pitch-based musical materials. After a presentation of the motivation for the project, theoretical background and related work, we present a set of research questions followed by a description of hardware and software approaches that address these questions. The paper concludes with a description of our plans to implement and embed these approaches in a robotic marimba player that will be used in workshops and concerts.

Keywords
Robotic Musicianship, perceptual modeling, algorithmic improvisation, human-machine interaction.

1. BACKGROUND AND MOTIVATION
The Robotic Marimba Player project aims at facilitating meaningful and inspiring melodic and harmonic interactions between humans and machine, leading to novel musical experiences and outcomes. The robot is designed to combine computational modeling of music perception, interaction, and improvisation, with the capacity to produce melodic and harmonic acoustic responses in physical and visual manners. We believe that real-time collaboration between human and robotic players can capitalize on the combination of their unique strengths to produce new and compelling music. Unlike computer- and speaker-based interactive music systems, a physical anthropomorphic robot can create acoustically rich and intuitive visual interactions with humans. The acoustic richness is achieved due to the complexities of real life systems, whereas in computer-generated audio, acoustic nuances require intricate design and are ultimately limited by the fidelity and orientation of speakers. Computer- and speaker-based interactive music systems are also hampered by their inanimate nature, which does not provide players and audiences with physical and visual cues that are essential for creating an expressive musical experience. For example, motion size often corresponds to loudness and gesture location often relates to pitch. These cues provide visual feedback, help performers anticipate and coordinate their playing, and create an engaging experience for the audience by providing a visual connection to the generated sound. The main motivation for this work, therefore, is to utilize robotics for the creation of an interactive musical system that would provide the visual cues and acoustic richness required for expressive and creative interaction with humans. In order to create intuitive and inspiring social collaboration, the robot is designed to analyze music based on computational models of human perception and to generate algorithmic responses that are humanly impossible (“listen like a human, improvise like a robot”). Extending previous work on a perceptual robotic drummer that was primarily focused on rhythm [38], the robot is designed to listen to, analyze, and play melodic and harmonic music. It is designed to infer musical meaning from live input based on a set of cognitive models of musical percepts such as melodic attraction [15, 20], tension [15, 19], and similarity [1, 11, 28]. The robot’s musical responses are designed to utilize mathematical constructs such as fractals [5, 12, 17, 29], cellular automata [4, 17] and genetic algorithms [13, 18], while its interaction schemes utilize synchronous and sequential operations [37], addressing aspects such as beat tracking [8, 27] and style adaptation [6, 22].

2. RELATED WORK
A number of research fields inform the design of the robotic marimba player, including musical robotics, human-robot interaction, and machine musicianship. In recent years, the field of musical robotics has received commercial, artistic, and academic interest, addressing a variety of musical instruments, including chordophones, aerophones, membranophones and idiophones. Several approaches have been explored for robotic stringed instruments utilizing solenoids, servomotors, and electro valves that slide and pick guitar strings in a variety of manners [14, 31]. Sophisticated control has been developed for robots such as the anthropomorphic robot flutist, which uses a complex mechanical imitation of human organs (including mechanical lungs, lips, fingers, and tongue) in an effort to accurately reproduce human subtleties [33]. Other examples for aerophone

robotic instruments include Toyota's Robotic Trumpeter [34] and the Rae's Autosax [25]. A number of percussive robots have been recently developed such as that ModBots – miniature modular instruments that can be affixed to virtually any structure [30], and the Thelxiepeia, which is designed to capture a wider timbral variety by using several actuators simultaneously [2]. Such research directions in musical robotics, however, focus on sound production and rarely address interactive and perceptual aspects of musicianship such as listening, analysis and improvisation. These research areas are explored by researchers in the field of machine musicianship [26]. Here, computational systems are developed to analyze, perform, and compose music with computers based on theoretical foundations in fields such as human-computer interaction, music theory, computer music, cognition, and artificial intelligence. One of the earliest research directions in this area is Score Following, where the computer synchronizes MIDI [7, 36], and recently audio [21] accompaniment to musical input from live soloists. Related efforts have been made to design computerized beat tracking applications that analyzes and match tempi and meter in real-time [8, 27] as well as improvisatory systems that responds to human input by manipulating a variety of musical parameters" [26], [16] [23]. The robotic marimba project attempts to combine the fields of musical robotics and the field of machine musicianship by developing a robot that utilizes perceptual modeling of high level musical aspects such as melodic attraction [15, 20], tension [15, 19], and similarity [1, 11, 28] as well as algorithmic improvisation based on computational constructs such as fractal algorithms [5, 17] cellular automata [4, 17], and genetic algorithms [13, 18].

3. PREVIOUS WORK BY AUTHORS

In previous work the authors have begun to explore the notion of robotic musicianship by developing an interactive robotic percussionist named Haile. Perceptually, Haile recognizes low-level musical aspects from drum playing such as note onset, pitch, and amplitude [24], as well as higher-level percepts such as rhythmic stability and similarity [9, 10]. The stability module, based on [9], rates the relationship between pairs of adjacent note durations according to their perceptual expectancy. Based on this model Haile listens to rhythmic phrases and modify them using desired stability and similarity parameters (see details in [39]). Mechanically, Haile controls two robotic arms; the right arm is designed to play fast notes, while the left arm is designed to produce larger and more visible motions, which can create louder sounds in comparison to the right arm. Unlike robotic drumming systems that allow hits at only a few discrete locations, Haile's arms can strike anywhere on a line between the center and the rim of the drum. Haile's interaction modules are based on the theory of interdependent group interaction in interconnected musical networks [37], using sequential and synchronous operations with centralized and decentralized control schemes. A number of interaction modes have been developed based on this theory such as Imitation, Stochastic Transformation, Perceptual Transformation, Simple Accompaniment, and Perceptual Accompaniment. In perceptual Accompaniment mode, for example, Haile plays simultaneously with human players while listening to and analyzing their input. It then creates local call-and-response interactions with different players based on the perception of rhythmic density and amplitude (see video clips here: http://www-static.cc.gatech.edu/~gilwein/Haile.htm). A user study conducted to evaluate Haile's human-robot interaction effectiveness [40] led to the identification of a number of directions for further research. Based on the user responses, we are currently extending our previous perceptual, mechanical, and interaction research to melodic and harmonic music. Findings from the study regarding the mechanics of the robotic percussionist emphasized the importance of anthropomorphic design for the creation of familiar and inspiring interactions with humans both rhythmically and melodically. The study also emphasized the need for more sophisticated perceptual analysis, larger and more visual movements and richer acoustic variety. Another important aspects evaluated in the user study addressed the prospect of creating novel human-robot musical experiences that cannot be generated by humans. Here, users' responses suggested the need for more richer algorithmic response and new playing techniques that will inspire humans to play and think about music in novel manners.

Figure 1: Haile - a perceptual robotic percussionist that listens and improvises with human drummers

4. RESEARCH QUESTIONS

A number of research questions guide the design of the robotic marimba player:

- Can we effectively implement computational schemes that model how humans represent and process melodic and harmonic structures in music? Can a robot use such models to infer high-level musical meaning form live musical input and respond in a musically intuitive manner?

- Can algorithmic models for musical improvisation create meaningful and inspiring musical responses? Can such algorithmic responses lead to novel socio-musical human-machine interaction and to music that cannot be created by humans?
- What is the role of physical, visual, and acoustic cues in multi-player musical interactions? Can a robot utilize physical properties to enrich musical interactions with humans?

5. HARDWARE DESIGN

In an effort to address these questions and to facilitate intuitive human-machine interaction through visual motions we decided to use an anthropomorphic design that can help convey the notion that the machine can listen and think. Findings from robotic percussionist user studies [40] stressed the importance of large visual motion in enabling humans to synchronize with and anticipate the robot's actions. The robotic marimba player, therefore, is designed to use several mallets that generate large and visible striking motion, both horizontally and vertically. Inspired by human playing techniques the design consists of four arms, each with three degrees of freedom, and a span of one octave (see Figure 2).

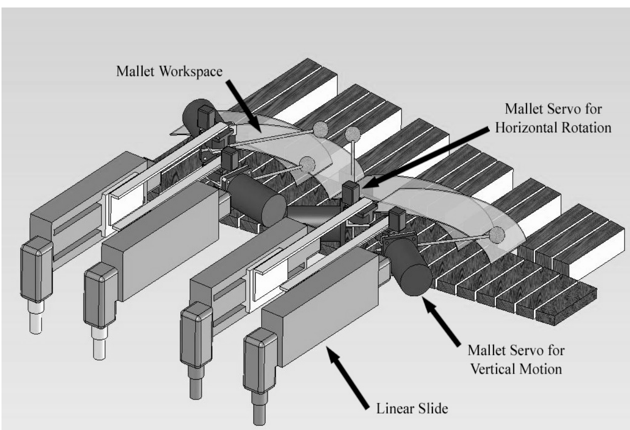

Figure 2: The design of the robotic marimba player consists of four arms, each with three degrees of freedom. A servomotor mounted at the end of each arm rotates mallets in a vertical plane to strike the bars. An additional servomotor and a linear slide assembly act together to position the mallets over an octave range.

The robotic arms are arranged in pairs with overlapping workspaces to allow various combinations of chords to be played. Four arms were chosen because marimba players typically hold four mallets (2 in each hand). However, due to the layout of the bars and necessary grips, human players must rotate their wrists in difficult angles to play certain chords, thus limiting their ability to quickly transition between such chords. Thanks to its four independent arms, the robot is only limited by the speed at which each mallet can move, although a certain amount of coordination between neighboring arms is required to avoid collisions in the shared workspace. The independent operation of each arm enables the robot to play sophisticated note combinations faster and more accurately than humans. One advantage humans have is their ability to shift their feet to play a larger range. This limitation is mitigated through a modular design that can be expanded to cover any number of octaves.

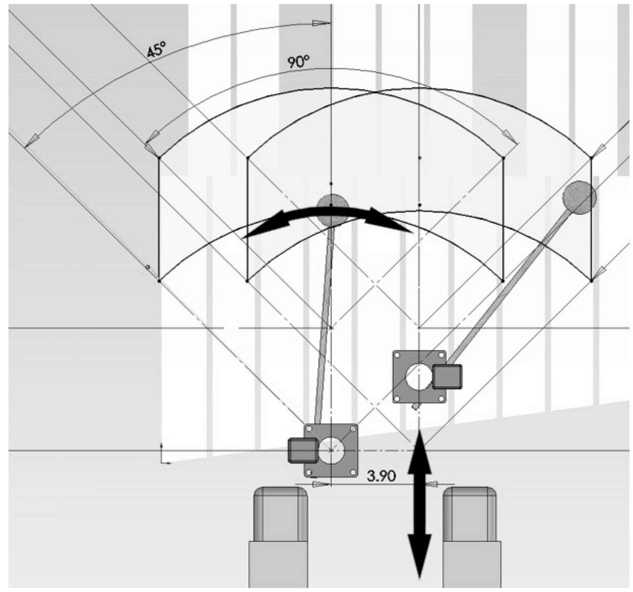

Figure 3: The pivot point of each mallet can be moved 14 cm towards and away from the marimba to play upper and lower keys, and the mallets can rotate laterally by 90 degrees to reach a one octave span.

As part of the design process of the robotic arms, several short recordings were made to obtain an approximate measure for human marimba playing speed. The recordings, made by Tom Sherwood from the Atlanta Symphony Orchestra, showed an average of about 12.5 notes per second (n/s) for chromatic scales, 10.7 n/s for major scales, and 8.7 n/s for arpeggios. Human speed depends on a number of factors, including the mallet weights, marimba bar spacing, and personal skill, therefore these measurements have been taken only as an approximate guide. Given fastest rate of 12.5 n/s using alternating hands, each robotic arm can operate at 6.25 n/s. However, since there are situations where only one hand is available to play a passage, or when extremely fast trills on one note are required, a maximum striking speed of 10Hz is designed for the mallet-striking servomotor. Based on mallet inertia of .0018kg-m2, the motor for each arm is designed to provide 2.5Nm of torque to achieve full stroke (0° to 50°) repetition at 10Hz, without accounting for motor inertia or friction. The other two axes, which locate the mallet over a desired note, affect the overall speed depending on the mallet's current location. The horizontal axis primarily determines the time needed to move between laterally adjacent notes, while the linear axis affects the time needed to switch between upper and lower notes (see corresponding bold arrows in Figure 3). Assuming a combined mallet and striker motor between 3 and 4 kg, the horizontal servo is modestly sized (~1Nm) and still cover an octave range in 0.25 seconds, (30 n/s). The linear slide, however, requires a large force to be able to match the 12.5 notes per second chromatic scale rate since it must move the weight of two actuators. As a compromise, it is sized to enable full up-down motion as shown in Figure 3 in 0.2 seconds, which only requires about 70 N.

Similarly to the robotic percussionist project, control over the robot is divided into several stages, covering high- to low-level aspects. At the highest-level, the computer issues note commands containing pitch, volume, and potentially a stroke style. This information is then passed to a path planning microprocessor where the current state of the machine is used to decide the most capable hand to generate the next note. This depends on the proximity of the mallets, the time since that mallet most recently played a note, and collision avoidance considerations. Once a mallet has been selected, trajectories for all three axes are calculated and sent as position set points to their respective position control loops. A special control scheme is utilized for the servomotor that is responsible for the vertical strike motion, depending on the style of stroke requested. For instance, if staccato notes are requested, the controller acts to keep the mallet from bouncing off the marimba to dampen the note quickly. Finally, the position controllers output the desired current set points that are utilized in the motor amplifiers.

6. SOFTWARE DESIGN

In order to facilitate intuitive melodic relationships with human players we are complementing our perceptual rhythmic modeling with a set of algorithmic models of melodic tension, attraction, and similarity as well as new models for sequential and synchronous musical human-robot interaction based on a theory of interconnected group collaboration [37]. In an effort to extend what is humanly possible, we are also working on a number of interactive improvisational algorithms based on fractal, cellular automata, and genetic algorithms. The research directions and theoretical frameworks for our approach are described below.

6.1 Perception

Our work in musical perception focuses on modeling and implementation of high-level musical percepts such as melodic attraction, tension and similarity. For low-level percepts such as note onset, amplitude, and pitch detection we are using known approaches for audio analysis [24] as well as input from MIDI instruments that generate discrete and accurate pitch, loudness and timing information. Based on a theoretical model for western tonal music we analyze monophonic melodic lines, addressing aspects such as melodic magnetism (the tendency of an unstable note to move up or down to the nearest tonal stable pitch), gravity (the tendency of an unstable note to move down to a lower note), and inertia (the tendency of musical patterns to continue in the same direction and attraction). The algorithm is designed to quantify the relative levels of these forces in live musical input and to use these parameters as coefficients in our improvisation engine (see section 6.2.3). Our melodic analysis is also informed by a quantitative approach for computing melodic attraction suggested by Lerdahl and Jackendoff [15], which returns a value for the attraction between pitches in a given tonality based on a table of anchoring strengths (see Table 1).

Table 1 – An anchoring strength table for computing the attraction between pitches, based on [15]

Anchoring Strength	The basic space with the fifth Omitted											
4	0											
3	0				4			7				
2	0		2		4	5		7		9		11
1	0	1	2	3	4	5	6	7	8	9	10	11

The table is based on the equation $delta(p1–p2) = (s2/s1)(1/n^2)$, where *delta*(p1–p2) is the level of attraction between pitch 1 and pitch 2, s2 and s1 are their respective anchoring strengths, and n is the distance between the two pitches in half steps. Table 1 shows that in the scale of C, b would be more attracted to c [*delta*(b–c) = 4/2 *1/1= 2] than d to c [*delta*(d–c)= 4/2 * 1/4 = 0.5.] Moreover, f would be more attracted to e [*delta*(f–e) = 3/2 * 1=1.5] than vice versa [*delta*(e–f) = 2/3 * 1 = 0.67]. These findings tend to confirm western listeners' intuitions, but the theory is regarded controversial due to its arbitrary theoretical choices (such as the value of n^2) and its narrow focus on local attractions. Eugene Narmour's approach for melodic attraction [19, 20] addresses some of these shortcomings. Narmour's main hypothesis is that any two successive pitches in western music imply a third pitch, while any pair of melodic pitches transmits separate intervallic and registral messages to the listener implying continuation when the interval is small (e.g. a major second going up followed by another major second going up) or reversal of direction and differentiation of interval size when it is large (e.g. a major sixth going up followed by a minor second going down). These concepts can be useful in codifying the analysis of melodic contour and how it affects the listeners' perception of tension. Similarly to Lehrdahl's and Jackendoff's melodic attraction rule, Narmour provides a quantitative measure of how pitch proximity and tonal anchoring [3] affect perceived tension, which we implement in our model. We also base our approach for melodic similarity on psychoacoustics studies that demonstrate the perceptual significance of melody contour for the perception of similarity [28]. In one case, it has been shown that novices' ability to retain melodic contour of a semi-known melody is much better than their ability to retaining the specific pitches [32]. Trehub demonstrated that contour can be perceived by infants as young as one year old, strengthening the assumption that this percept is well ingrained in human cognition as a similarity perception construct [35].

6.2 Interaction

In order to utilize the analysis and rating of high-level melodic percepts we have developed interaction schemes based on sequential, continuous, centralized, and decentralized interconnections, in both symmetric and asymmetrical network topologies. Based on this framework, we are currently working on a number of interaction models that are inspired by human-human interaction such as beat detection and style adaptation. Creating an effective synchronization and tempo tracking is one of the important skills needed for humans to coordinate their musical actions with each other. Effective beat tracking would not only enable timely accompaniment, but would also facilitate interaction schemes where the robot augments source material to rhythmically fit with current human output. Our beat detection approach, based on [8, 27], has already provided successful results – see preliminary examples with the robotic percussionist at http://coa.gatech.edu/~gil/Beatdetection.mov. We are designing improvements to this algorithm that include anticipation factors based on measurement of acceleration and deceleration, as well as periodic sequences of listening and playing that can help stabilize the beat detection model and appropriate it for human-machine interaction. Another approach under consideration is based on the fact that when humans improvise collaboratively, they tend to exchange musical motifs and gestures, copying and extending each other's ideas and styles. In future work, using learning algorithms such an augmented Markov models for musical styles [22] the marimba robotic player will be designed to learn humans musicians' style in terms of their high-level perceptual input and will use this analysis for its response.

6.3 Improvisation

Our work on machine improvisation aims to extend human-human interaction by capitalizing on computational capabilities such as processing power, the ability to perform sophisticated mathematical transformations, robust long-tem memory, and the ability to play accurately without practice. We are not only interested in operatins that can be generated faster and more accurately than humans (such as inverting, retrograding, looping or shifting musical materials in time), but also in novel musical responses that cannot be replicated even by the most skilled human musician. A promising approach in the regard is the utilization of cellular automata, fractals, and genetic algorithms. The graphical artwork generated by fractal algorithms such as the Mandelbrot set, for example, contains structural features that can also be found in music, such as similarity across different scales and repeated patterns with nuanced differences. It is unlikely for humans to create such patterns due to the number and speed of mathematical operations needed and the recursive nature of the function. However, while the mathematical algorithm is ultimately responsible for the patterns formed, human decisions are necessary to transform a string of numbers into art with aesthetic merit. For fractal graphics, this often involves color-mapping, coefficient range decisions, and choosing output section regions. In music too, the aesthetics of the result depends on human decisions. Several previous works in fractal music addressed the analysis and composing of self contained musical structures in a non-interactive manner [5, 12, 17, 29]**.** Our approach, on the other hand, explores schemes that allow for live interaction between humans and robots, injecting human expression into the robotic fractal algorithms. For example, within a fractal set, there are large regions of asymptotic behavior where iterations quickly converge towards zero or infinity. On the edges between those regions, there is a chaotic region containing oscillating patterns that bear interesting behaviors that can be applied to music. Some patterns continue on indefinitely without change, others slowly oscillate between states, and still others unpredictably cross a threshold that causes them to quickly converge towards infinity or zero.

To convert this behavior to music, our design maps short rhythmic phrases to regions in a fractal's output space. Those phrases are then played every time the fractal algorithm enters a predefined region. By mapping the characteristics of an accompanying human to coefficients in the fractal algorithm, the human player can affect its direction. In our preliminary design, the stability percept is mapped to fractal stability, so that the outcome is more likely to become unstable when the human plays in an unstable manner. The recursive nature of some fractals can also lead to sophisticated and unpredictable results when a function is repeatedly applied to a musical phrase. The robot can, therefore, serve as an inherently unpredictable source of ideas, while human players can restrain and shape the musical outcome by controlling coefficients in the function. Other examples for algorithms that bear great promise for the creation of aesthetical musical patterns and structures include cellular automata and genetic algorithms, which utilize simple decentralized rules that can create patterns with varied levels of stability. These can be mapped to rhythmic and melodic output, while humans control the starting position and the rules of operation in real time. Genetic algorithms, on the other hand, can operate on large "populations" of human- and artificially- generated musical motifs and use techniques inspired by evolutionary biology such as mutation, selection, and crossover to shape the musical output based on a fitness function controlled by similarity analysis from human input

7. FUTURE WORK

The development of the perceptual and improvisational robotic marimba player is work in progress. In future work we indented to continue the development of the computational ideas described above such as harmonic perception, beat detection, style adaptation, cellular automata and genetic algorithms. We than plan to implement these modules in the robotic hardware based on the design described in this paper. We also intend to conduct user studies leading to workshops and concerts with the robot. The user studies will be designed to evaluate the effectiveness of our approaches in perception, mechanics and interaction design. Perceptually, we will compare between the results produced by the perceptual models and human listeners' subjective perception of melodic attraction, tension and similarity. These studies will utilize qualitative methods similarly to our previous studies on rhythmic stability and similarity [39]. In order to assess the mechanical musical abilities of the robot we will use both quantitative and qualitative evaluation methods. As part of the quantitative evaluation we will measure the robot's time and error of navigation, pitch and dynamic range, resolution accuracy, and latency in comparison to human subjects in different levels of skills. We also plan to design a number of music specific tasks for comparison such as timbre control, musical gestures (glissandi, trills, vibrato, grace notes) musical phrase (control speeds and articulations, scales, arpeggios), as well as task combinations. The qualitative section of the study will include a user survey, where subjects will be asked about the strengths and weaknesses of different aspects in the mechanical design. In the interaction and improvisation evaluation sessions, subjects will be asked to play and improvise with other humans, with a software synthesizer version of the interaction and improvisation models (using speakers as output), as well as with the same models as implemented in the robot. Subjects will then be asked to assess the effectiveness of the interaction in terms of visual cues, synchronization, anticipation, acoustic richness, style adaptability, and beat tracking. Findings from the user studies will be used to improve the hardware and software modules, towards workshops and concerts that we plan to conduct with educational and artistic organizations.

8. ACKNOWLEDGMENTS

We would like to thank the Music Department the GVU center at Georgia Tech for supporting this project.

9. REFERENCE

[1] Arom, S. 1985. De L'ecoute a l'analyse des musiques centrafricaines. *Analyse Musicale*, Vol. 1, p. 35-59.

[2] Baginsky, N.A. 2004. *The Three Sirens: A Self Learning Robotic Rock Band*. http://www.the-three-sirens.info.

[3] Bharucha, J. 1996. Melodic Anchoring. *Music Perception*, Vol. 13, p. 282-400.

[4] Bilotta, E. 2002. Synthetic Harmonies: An Approach to Musical Semiosis by Means of Cellular Automata. *Leonardo*, Vol. 35(2), p. 153-159.

[5] Chapel, R., *Realtime Algorithmic Music Systems From Fractals and Chaotic Functions: Toward an Active Musical Instrument*, in *Department of Technology*. Doctorate Thesis, 2003, Universitat Pompeu Fabra: Barcelona.

[6] Cope, D. 1991. *Computers and Music Style*. Oxford: A-R Editions.

[7] Dannenberg, R. An On-line Algorithm for Real-Time Accompaniment. in *International Computer Music Conference* 1984. Paris, France.

[8] Davies, M. and M. Plumbley. Beat Tracking With A Two State Model. in *IEEE International Conference on Acoustics, Speech and Signal Processing*. 2005. Philadelphia, USA.

[9] Desain, P. and H.J. Honing. Rhythmic stability as explanation of category size. in *of the International Conference on Music Perception and Cognition* 2002. Sydney.

[10] Dibben, N. 1999. The Perception of Structural Stability in Atonal Music: The Influence of Salience, Stability, Horizontal Motion, Pitch Commonality, and Dissonance. *Music Perception*, Vol. 16(3), p. 265-295.

[11] Hewlett, B. and E. Selfreidge-Field 1998. *Melodic Similarity - Concepts, Procedures, and Applications*. Computing in Musicology. Cambridge, MA: MIT Press.

[12] Hsu, K. and A. Hsu 1990. Fractal Geometry of Music. *Proceedings of the National Academy of Sciences,*, Vol. 87, p. 938-941.

[13] Jacob, B. Composing with Genetic Algorithms. in *International Computer Music Conference*. 1995. Bannf Alberta.

[14] Jordà, S. Afasia: The Ultimate Homeric One-man multimedia- band. in *The International Conference on New Interfaces for Musical Expression*. 2002.

[15] Lerdahl, F. and R. Jackendoff 1983. *A Generative Theory of Tonal Music*. Cambridge, Massachusetts: MIT Press.

[16] Lewis, G. 2000. Too Many Notes: Computers, Complexity and Culture in Voyager. *Leonardo Music Journal* Vol. 10 p. 33-39.

[17] McAlpine, K., E. Miranda, and S. Hoggar 1999. Making Music with Algorithms: A Case-Study System. *Computer Music Journal*, Vol. 23(12), p. 19-30.

[18] Moroni, A. 2000. Vox Populi: An Interactive Evolutionary System for Algorithmic Music Composition. *Leonardo*, Vol. 10, p. 49-54.

[19] Narmour, E. 1990. *The Analysis and Cognition of Basic Melodic Structures*. Chicago: University of Chicago Press.

[20] Narmour, E. 1992. *The Analysis and Cognition of Melodic Complexity: The Implication-Realization Model*. Chicago: University of Chicago Press.

[21] Orio, N., S. Lemouton, and D. Schwarz. Score Following: State of the Art and New Developments. in *International Conference on New Interfaces for Musical Expression*. 2003. Montreal, Canada.

[22] Pachet, F. 2003. The Continuator: Musical Interaction With Style. *Journal of New Music Research*, Vol. 32(3), p. 333-341.

[23] Pachet, F. The continuator: Musical interaction with style. in *International Computer Music Conference*. 2002. Göteborg, Sweden.

[24] Puckette, M. Real-time Audio Analysis Tools for Pd and MSP. in *International Computer Music Conference*. 1998.

[25] Rae, G.W. 2005. *Robotic Instruments*. http://logosfoundation.org/instrum_gwr/automatons.html

[26] Rowe, R. 1992. *Interactive Music Systems: Machine Listening and Composing*. Cambridge, MA: MIT Press.

[27] Scheirer, E. 1998. Tempo and beat analysis of acoustic musical signals. *Journal of the Acoustical Society of America*, Vol. 103(1), p. 588-601.

[28] Schmuckler, M.A. 1999. Testing Models of Melodic Contour Similarity. *Music Perception*, Vol. 16(3), p. 295-326.

[29] Schroeder, M. 2002. Is there such a thing as fractal music? *Nature*, Vol. 325, p. 765-766.

[30] Singer, E., J. Feddersen, R. C., and B. Bowen. LEMUR's Musical Robots. in *International Conference on New Interfaces for Musical Expression*. 2004. Hamamatsu, Japan.

[31] Singer, E., K. Larke, and D. Bianciardi. LEMUR GuitarBot: MIDI Robotic String Instrument. in *International Conference on New Interfaces for Musical Expression*. 2003. Montreal, Canada.

[32] Sloboda, J. 1985. *The Musical Mind*. Oxford: Clarendon Press.

[33] Takanishi, A. and M. Maeda. Development of Anthropomorphic Flutist Robot WF-3RIV. in *Proceedings of International Computer Music Conference*. 1998.

[34] Toyota 2004. *Trumpet Robot*. http://www.toyota.co.jp/en/special/robot.

[35] Trehub, S.E., D. Bull, and L.A. Thorpe 1984. Infants' perception of melodies: The role of melodic contour. *Child Development*, Vol. 55, p. 821-830.

[36] Vercoe, B. "The Synthetic Performer in the Context of Live Performance. in *International Computer Music Conference*. 1984. Paris, France.

[37] Weinberg, G. 2005. Interconnected Musical Networks - Toward a Theoretical Framework. *Computer Music Journal*, Vol. 29(2), p. 23-39.

[38] Weinberg, G. and S. Driscoll. Robot-Human Interaction with an Anthropomorphic Percussionist. in *International ACM Computer Human Interaction Conference* 2006. Montréal, Canada.

[39] Weinberg, G., S. Driscoll, and M. Parry. Musical Interactions with a Perceptual Robotic Percussionist. in *of IEEE International Workshop on Robot and Human Interactive Communication* 2005. Nashville, TN.

[40] Weinberg G. and S. Driscoll 2006. Towards Robotic Musicianship. *Computer Music Journal*, Vol. 30(4), p. 28-45.