We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Robotic Musicianship – Musical Interactions Between Humans and Machines

Gil Weinberg
Georgia Institute of Technology
USA

1. Introduction

The Robotic Musicianship project aims to facilitate meaningful musical interactions between humans and machines, leading to novel musical experiences and outcomes. The project combines computational modelling of music perception, interaction, and improvisation, with the capacity to produce acoustic responses in physical and visual manners. The motivation for this work is based on the hypothesis that real-time collaboration between human and robotic players can capitalize on the combination of their unique strengths to produce new and compelling music. Our goal is to combine human qualities such as musical expression and emotions with robotic traits such as powerful processing, the ability to perform sophisticated mathematical transformations, robust long-term memory, and the capacity to play accurately without practice. A similar musical interaction can be achieved with software applications that do not involve mechanical operations. However, software-based interactive music systems are hampered by their inanimate nature, which does not provide players and audiences with physical and visual cues that are essential for creating expressive musical interactions. 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 musical experience by providing a visual connection to the generated sound. Software based interactive music systems are also limited by the electronic reproduction and amplification of sound through speakers, which cannot fully capture the richness of acoustic sound. Unlike these systems, the anthropomorphic musical robot we developed, named Haile, is designed to create acoustically rich interactions with humans. The acoustic richness is achieved due to the complexities of real life systems, as opposed to digital audio nuances that require intricate design and that are limited by the fidelity and orientation of speakers. In order to create intuitive as well as inspiring social collaboration with humans, Haile is designed to analyze music based on computational models of human perception and to generate algorithmic responses that are unlikely to be played by humans ("listen like a human, improvise like a machine"). It is designed to serve as a test-bed for novel forms of musical human-machine interactions, bringing perceptual aspects of computer music into the physical world both visually and acoustically. We believe that this approach can lead to new musical experiences, and to new music, which cannot be conceived by traditional means.

Source: Human-Robot Interaction, Book edited by Nilanjan Sarkar,
ISBN 978-3-902613-13-4, pp.522, September 2007, Itech Education and Publishing, Vienna, Austria

2. Related Work

Two main research areas inform our effort to develop robotic musicianship – *musical robotics*, which focuses on the construction of automated mechanical sound generators and *machine musicianship*, which centres on computer models of music theory, composition, perception, and interaction. Early work on musical robotics focused on mechanical keyboard instruments such as the Pianista by French inventor Fourneax (see a comprehensive historic review of musical robots in (Kapur 2005)). In recent years, the field has received commercial, artistic, and academic interest, expanding to anthropomorphic designs as well as robotic musical instruments, including chordophones, aerophones, membranophones and idiophones. Several approaches have been recently explored for robotic stringed instruments. GuitarBot (Singer et al. 2004), for example, is a mechanical guitar operated by a set of DC servomotors driving a belt with multiple picks playing four strings. The pick position, controlled by a photosensor and a “clapper” solenoid, is used as a damper. Jordà’s Electric Guitar Robot (Jordà 2002), on the other hand, has six strings, that can be plucked by twelve picks, driven by an electro-valve hammer-finger. Current approaches for mechanical guitars, however, are not designed to explore the full range of sonic variety through string techniques such as bouncing, bowing, strumming, scratching, or rubbing. Other attempts have been made to develop expressive wind instrument robots. The Anthropomorphic Flute Robot (Chida, Okuma et al. 2004), uses a complex mechanical imitation of human organs in an effort to accurately reproduce human flute playing. The elaborate apparatus includes robotic lungs, neck, lips, fingers, and tongue. Other examples for aerophone robotic instruments are Toyota’s Robotic Trumpeter (Toyota 2007) and the Rae’s Autosax (Rae 2005), which are programmed to follow deterministic rules. More varied work has been done on robotic percussionists, both for idiophone and membranophone instruments. The ModBots (Singer et al. 2004), for example, are miniature modular instruments designed to affix to virtually any structure. Each ModBot consists of only one electromechanical actuator (a rotary motor or a linear solenoid), which responds to varying degrees of supply voltage regulated by a microcontroller. A more elaborated mechanism by Singer is utilized in the TibetBots, which consist of six robotic arms that strike three Tibetan singing bowls. Here, an effort was made to capture a wider timbral variety by using two robotic arms (controlled by solenoids) for each bowl to produce a richer set of sounds. Another approach for broadening timbre and pitch versatility is taken by the Thelxiepeia (Baginsky 2004). The instrument consists of a mechanical drumstick and a motorized mechanism to rotate the drum circumference, which can lead to the production of a range of pitches. Other robotic instruments which influenced our work were developed by Trimpin (Trimpin 2000), Rae (Rae 2005), and Van Doressen (Dorssen 2006).

The second research area that informs our work is machine musicianship. Here, researchers design and develop computer systems that analyze, perform, and compose music based on theoretical foundations in fields such as music theory, computer music, cognition, artificial intelligence and human-computer-interaction (Rowe 1992). One of the earliest research directions in this field is the “Score Follower”, in which the computer tracks a live soloist and synchronizes MIDI (Dannenberg 1984) (Vercoe 1984), and recently audio (Orio, Lemouton et al. 2003), accompaniment to the musical input. The classic score following approach focuses on matching predetermined musical events to real-time input. A more improvisatory approach is taken by systems such as Voyager (Lewis 2000) and Cypher (Rowe 1992). Here the software analyzes musical input in real time and generates musical

responses by manipulating a variety of parameters such as melody, harmony, rhythm, timbre, and orchestration. David Cope has taken a non-real time approach in his system for analyzing composers' styles based on MIDI renditions of their compositions (Cope 1996). Cope's algorithm learns the style of a given composer by modelling aspects such as expectation, memory, and musical intent. It can then generate new compositions with stylistic similarities to the originals. The "Continuator" system, on the other hand, takes a real-time approach for learning the improvisation style of musicians as they play polyphonic MIDI instruments (Pachet 2003). The application uses Hidden Markov Models to learn and analyze the input and continues the improvisation in the style of the human performer.

Particularly note-worthy research field in machine musicianship is computational modelling of music perception in which, researches develop cognitive and computational models of low- and high-level musical percepts. Lower level cognitive modelling address percepts from note onset detection to pitch and beat detection, using audio sources (Puckette 1998) (Scheirer 1998) (Foote and Uchihashi 2001) as well as MIDI (Winkler 2001). Higher-level rhythmic percepts include more subjective concepts such as rhythmic stability, melodic similarity and attraction. Desain and Honing's model of rhythmic stability is based on the relationship between pairs of adjacent note durations (Desain and Honing 2002); Tanguiane counts the number of coincident onset in an effort to model rhythmic similarity of different audio signals (Tanguiane 1993); Smith utilizes dynamic time warping techniques to retrieve similar melodies from a folk song database (Smith et al. 1998); and Lerdahl and Jackendoff calculate the melodic attraction between pitches in a given tonality based on a table of anchoring strengths (Lerdahl & Jackendoff 1983). Informed by such approaches for perceptual modelling, our robot is designed to respond with algorithmically generated musical outcomes using a novel approach for computational composition and improvisation. This aspect of the system is based on theoretical approaches for musical improvisation and interaction (Pressing 1994), (Johnson-Laird 2002), as well as practical efforts using methods such as genetic algorithms. GenJam, for example, is an interactive computer system that improvises over a set of jazz tunes using an initial phrase population that is generated stochastically (Biles 1994). GenJam's fitness function is based on human input, where in every generation the user determines which phrases remain in the population. Other systems use methods such as real-time fitness criteria (Moroni 2000) or human feedback for training a neural network-based fitness function (Tokui & Iba 2000). In the second phase of the robotic musicianship project we developed an improvisatory genetic algorithm that combines human aesthetics and perception with algorithmic "gene mixing" improvisation.

3. Research questions

A number of research questions guide our effort to create intuitive and inspiring musical human-robot interactions and to establish the concept of robotic musicianship:

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

Below we describe our efforts to address these research questions through physical and mechanical design (section 4), rhythmic and melodic applications (sections 5), user studies (section 6), and a number of directions for future work (section 7).

4. Physical and Mechanical Design

Figure 1. Haile's design

In order to support familiar and expressive interactions with human players, Haile's design is anthropomorphic, utilizing two percussive arms that can move to different locations and strike with varying velocities. The first prototype was designed to play a Native American Pow Wow drum – a unique multi player instrument that supports the collaborative nature of the project. For pitch-oriented applications, the robot was later adjusted to play a one-octave xylophone. In order to match the aesthetics of these musical instruments, we chose to construct the robot from wood. The wooden parts were made using a CnC wood cutting machine and constructed from several layers of plywood glued together. Metal joints were designed to allow shoulder and elbow movement as well as leg adjustability for different instrument heights. While attempting to create an organic look for the robot, it was also important that the technology was not completely hidden, so that co-players could see and understand the robot's operation. We therefore left the mechanical apparatuses uncovered and embedded a number of LEDs on Haile's body, providing an additional representation of the mechanical actions (See Figure 1).

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 that produce louder sounds. Both arms can adjust the strikes sound in two manners: different pitches are achieved by striking the instruments in different locations, and volume is adjusted by hitting with varying velocities. To move to different vertical positions, each arm employs a linear slide, a

belt, a pulley system, and a potentiometer to provide feedback (see Figure 2). Unlike robotic drumming systems that allow hits at only a few discrete locations, Haile's arms move continuously over a distance of 10 inches (movement timing is 250 ms. from end to end). The right arm's striking mechanism is loosely based on a piano hammer action and consists of a solenoid driven device and a return spring (see Figure 3). The arm strikes at a maximum speed of 15 Hz, faster than the left arm's maximum speed of 11 Hz. However, the right arm cannot generate a wide dynamic range or provide easily noticeable visual cues, which limits Haile's expression and interaction potential. The left arm was designed to address these shortcomings, using larger visual movements, and a more powerful and sophisticated hitting mechanism. Whereas the striking component of the right arm is about the size of a finger and can only move 2.5 inches vertically (see figure 4), the entire left forearm takes part in the striking motion and can move up and down eight inches. A linear motor and an encoder located at the left elbow are used to provide sufficient force and control for the larger mass and motions (see Figure 5).

Figure 2. The right arm slider mechanisms

Figure 3. The right arm striking mechanism

Max/MSP, a graphical programming environment (Cycling74 2007), was used for high-level musical programming in an effort to make the project accessible to composers, performers, and students. The first right arm prototype incorporated the USB based Teleo System (MakingThings 2005) as the main interface between Max/MSP and Haile's sensors and motors. Low-level control of the solenoid-based right arm's position was computed within Max/MSP, which required a continuous feed of position updates to the computer. This consumed much of the communication bandwidth as well as processor time on the main computer. The final two-arm mechanism utilizes multiple onboard microprocessors for local low-level control as well as Ethernet communication with the main computer. The new

system, therefore, facilitates faster and more sophisticated control (2ms control loop) and requires only low bandwidth communications with the operating computer. Each arm is locally controlled by an 18F452 PIC microprocessor, both of which receive RS232 communications from a Modtronix Ethernet board (SBC68EC). The Ethernet board receives 3 byte packets from the computer, a control byte and two data bytes. The protocol utilizes an address bit in the control byte to send the information to the appropriate arm processor. The two data bytes typically contain position and velocity set points for each hit, but can also be used to update the control parameters.

Figure 4. Haile's right arm design

Figure 5. Haile's left arm design

Two onboard PIC microprocessors are responsible for controlling the arms' sliding and hitting mechanisms, ensuring that the impacts occur at the requested position and velocity. In order to allow enough time for the arms to move to the correct location and execute the strokes, a 300 ms. delay line is implemented between signal reception and impact. It has been shown that rhythmic errors of only 5 ms are detectable by average listeners (Coren, et al,

2003), therefore, it was important to ensure that this delay remained accurate and constant regardless of different hit velocities, allowing to easily compensate for it in the higher-level interaction application. Both arms store incoming hit commands in a First-In-First-Out queue, moving towards the location of a new note immediately after each hit. Due to its short vertical hitting range, the solenoid driven right arm allows for fairly consistent stroke time. We, therefore, implemented the 300 ms delay as a constant for this arm. The left arm, on the other hand, undergoes much larger movements, which requires complex feedback control to ensure that impact occurs at the right time, regardless of hits velocity. While waiting for incoming notes, the left arm remains about one inch above the surface of the instrument. When a new note is received, the arm is raised to a height proportional to the loudness of the hit. After a delay determined by the desired velocity and elevation, the arm descends towards the instrument under velocity control. After impact, the arm returns back to its standby position above the instrument. Extremely fast notes utilize a slightly different control mode that makes use of the bounce of the arm in preparation for the next hit. This mechanism allows the left arm to control a wide dynamic range and provides performers and viewers with anticipatory and real-time visual cues, enhancing expression and enriching the interaction representation.

5. Applications

5.1 Phase one – Rhythmic Interaction

The first phase of the project aimed at facilitating rhythmic collaboration between human drummers and Haile, addressing aspects such as rhythmic perception, improvisation, and interaction. In perception, we developed models for low- and high-level rhythmic percepts, from hit onset, amplitude, and pitch detection, through beat and density analysis, to rhythmic stability and similarity perception. For hit onset and amplitude detection we adjusted the Max/MSP *bonk~* object (Puckette 1998) to address the unique character of the Pow Wow drum- a multi player Native American percussion instrument, which was chosen for the project due to its collaborative nature. *Bonk~* provides effective onset attack detection but its frequency bands analysis is insufficient for accurate pitch detection due to the Pow Wow drum's low frequency and long reverberating sounds. Since *bonk~* is hard-coded with a 256 point analysis window, the lowest frequency it can analyze is 172Hz – too high for the Pow Wow drum which has a natural frequency of about 60 Hz. Moreover, onset detection is complicated when high frequency hits are masked by the long decay of the previous low strikes (see Figure 6). To address these issues, we wrote a Max/MSP external object that used 2048 points FFT to determine both the magnitude of lower frequency bins and the change in those magnitudes between successive analysis frames. By taking into account the spectral changes in addition to the magnitudes, Haile could better determine whether energy in a particular frequency band came from a current hit or from previous ones (See Figure 6).

Other relatively low-level perceptual modules that were developed were beat detection, where domain detection was followed by autocorrelation of tempo and phase (Davies and Plumbley 2005) and density detection, where we looked at the number of note onsets per time unit to represent the density of the rhythmic structure. We also implement a number of higher-level rhythmic analysis modules for percepts such as rhythmic stability, based on (Desain and Honing 2002), and similarity based on (Tanguiane 1993). The stability model calculates the relationship between pairs of adjacent note durations, rated according to their perceptual expectancy based on three main criteria: perfect integer relationships are favoured,

ratios have inherent expectancies (i.e., 1:2 is favoured to 1:3 and 3:1 is favoured to 1:3), and durations of 0.6 seconds are preferred. The expectancy function can be computed as:

$$E_b(A,B)=\int_0^r(\text{round}(r)-r)\times|2(r-\text{floor}(r)-0.5)|^p\times\text{round}(r)^d\,dr$$

where A and B are the durations of the two neighbouring notes, $r=\max(A/B,B/A)$ represents the (near) integer relationship between note durations, p controls the shape of the peaks, and d is negative and affects the decay rate as the ratios increases. This function is symmetric around $r=1$ when the total duration is fixed (see Figure 7a). Generally, the expectancy function favours small near-integer ratios and becomes asymmetric when the total duration varies, exhibiting the bias toward the 600 ms. interval (see Figure 7b).

Figure 6. Magnitude plots from a 60Hz, 300Hz, and 5kHz frequency band over several low and high-pitched hits showing the relatively slow decay of the low-pitched hits

Figure 7. Basic expectancy of intervals A and 1-A (a) and 0.3 and B (b), reproduced from (Desain and Honing 2002)

Our similarity rating is derived from Tanguiane’ binary representation, where two rhythms are first quantized, and then given a score based on the number of note onset overlaps and near-overlaps. In order to support real-time interaction with human players, we developed two Max/MSP externals that analyzed and generated rhythms based on these stability and similarity models. These externals were embedded in a live interaction module that read measure-length rhythmic phrases and modified them based on desired stability and similarity parameters. Both parameters varied between 0 and 1 and were used together to select an appropriate rhythm from a database of pre-analyzed rhythms. A stability rating of 1 indicated the most stable rhythm in the database, 0.5 equated to the stability of the input rhythm, and 0 to the least stable rhythm. The similarity parameter determined the relative contribution of similarity and stability.

The main challenge in designing the rhythmic interaction with Haile was to implement our perceptual modules in a manner that would lead to an inspiring human-machine collaboration. The approach we took to address this challenge was based on a theory of interdependent group interaction in interconnected musical networks (Weinberg 2005). At the core of this theory is a categorization of collaborative musical interactions in networks of artificial and live musicians based on sequential and synchronous operations with centralized and decentralized control schemes. For example, in sequential decentralized interactions, players create their musical materials with no influence from a central system or other players and can then interact with the algorithmic response in a sequential manner (see Figure 8). In a synchronous centralized network topology, on the other hand, players modify and manipulate their peers’ music in real-time, interacting through a computerized hub that performs analysis and executes generative functions (see Figure 9). More sophisticated schemes of interaction can be designed by combining centralized, decentralized, synchronous, and sequential interactions in different directions, and by embedding weighted gates of influence among participants (see Figure 10).

Figure 8. A model of sequential decentralized interaction. Musical actions are taken in succession without synchronous input from other participants, and with no central system to coordinate the interaction

Figure 9. A model of synchronous centralized interaction. Human and machine players are taking musical actions simultaneously, and interact through a computerized hub that interpret and analyze the input data

Figure 10. A combination of centralized, decentralized, synchronous and sequential musical actions in an asymmetric topology with weighted gates of influence. Based on (Weinberg 2005)

Based on these ideas, we developed six interaction modes for Haile: Imitation, Stochastic Transformation, Perceptual Transformation, Beat Detection, Simple Accompaniment, and Perceptual Accompaniment. These interaction modes utilize different perceptual modules and can be embedded in different combinations in interactive compositions and educational activities. In the first mode, Imitation, Haile merely repeats what it hears based on its low-level onset, pitch, and amplitude perception modules. Players can play a rhythm and after a couple of seconds of inactivity Haile imitates it in a sequential call-and-response manner. Haile uses one of its arms to play lower pitches close to the drumhead centre and the other arm to play higher pitches close to the rim. In the second mode, Stochastic Transformation, Haile improvises in a call-and-response manner based on players' input. Here, the robot stochastically divides, multiplies, or skips certain beats in the input rhythm, creating variations of users' rhythmic motifs while keeping their original feel. Different transformation coefficients can be adjusted manually or automated to control the level of similarity between users' motifs and Haile's responses. In the Perceptual Transformation mode, Haile analyzes the stability level of users' rhythms, and responds by choosing and playing other rhythms that have similar levels of stability to the original input. In this mode Haile automatically responds after a specified phrase length. Imitation, Stochastic Transformation, and Perceptual Transformation are all sequential interaction modes that form decentralized call-and-response routines between human players and the robot. Beat Detection and Simple Accompaniment modes, on the other hand, allow synchronous interaction where humans play simultaneously with Haile. In Beat Detection mode, Haile tracks the tempo and beat of the input rhythm using complex domain detection function and autocorrelation, which leads to continuously refined assumptions of tempo and phase. A simpler, yet effective, synchronous interaction mode is Simple Accompaniment, where Haile plays pre-recorded MIDI files so that players can interact with it by entering their own rhythms or by modifying elements such as drumhead pressure to modulate and transform Haile's timbres in real-time. This synchronous centralized mode allows composers to feature their structured compositions in a manner that is not susceptible to algorithmic transformation or significant user input. The Simple Accompaniment mode is also useful for sections of synchronized unisons where human players and Haile play together. Perhaps the most advanced mode of interaction is the Perceptual Accompaniment mode, which combines synchronous, sequential, centralized and decentralized operations. Here, 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 its perceptual analysis. In this mode we utilize the amplitude and density perceptual modules that are described above. While Haile plays short looped sequences (captured during the Imitation and Stochastic Transformation modes) it also listens to and analyzes the amplitude and density curves of human playing. It then modifies its looped sequence, based on the

amplitude and density coefficients of the human players. When the rhythmic input from the human players is dense, Haile plays sparsely, providing only the strong beats and allowing humans to perform denser solos. When humans play sparsely, on the other hand, Haile improvises using dense rhythms that are based on stochastic and perceptual transformations. Haile also responds in direct relationship to the amplitude of human players so that the louder humans play, the stronger Haile plays to accommodate the human dynamics, and vice versa (see a video excerpts of some of the interaction modes at <http://www.cc.gatech.edu/~gilwein/Haile.htm>.)

Figure 11. The composition Jam'aa, as performed at the RoboRave Festival in Odense, Denmark

As a creative outcome for these rhythmic applications, two compositions were written for the system, each utilized a different set of perceptual and interaction modules. The first composition, titled Pow, was written for one or two human players and a one-armed robotic percussionist. It served as test case for Haile's early mechanical, perceptual, and interaction modules. The second composition, titled Jam'aa ("gathering" in Arabic), builds on the unique communal nature of the Middle Eastern percussion ensemble, attempting to enrich its improvisational nature, call-and-response routines, and virtuosic solos with algorithmic transformation and human-robotic interactions. Here, the sonic variety of the piece was enriched by using two robotic arms and by including other percussive instruments such as darbukas (goblet shaped middle-eastern hand drum), djumbes, and tambourines. In Jam'aa Haile listens to audio input via directional microphones installed inside two darbuka drums played by humans. In some sections of the piece the left arm merely provides the beat while in other sections it participates in the algorithmic interaction. Jam'aa utilizes interaction modes that were not included in Pow, such as perceptual transformation and perceptual

accompaniment. We also developed a new response algorithm for Jam'aa titled "morphing", where Haile combines elements from two or more of the motifs played by humans, based on a number of integration functions. Jam'aa, was commissioned by Hamaabada Art Centre In Jerusalem, and later performed in invited and juried concerts in France, Germany, Denmark, and the United States (see a video excerpts from Jam'aa at <http://coa.gatech.edu/~gil/RoboraveShort.mov>)

5.2 Phase Two – Melodic Interaction

As part of our effort to expand the exploration of robotic musicianship into pitch and melody, Haile was adapted to play a pitch-based mallet instrument. The one-octave xylophone we built for this purpose was designed to fit Haile's mechanical design – the left arm covered a range of 5 keys while the right arm, whose vertical range was extendable, covered a range of 7 keys. The different mechanisms driving each arm led to unique timbres, as notes played by the solenoid-driven arm sound different than those played by the linear-motor based arm. Since the robot could play only one octave, the algorithmic responses were filtered by pitch class.

Figure 12. Haile's two robotic arms cover a range of one octave – from middle G to treble G

Following the guideline "listen like a human, improvise like a machine", we decided to implement a perceptual model of melodic similarity ("listen like a human") as the fit function of a genetic algorithm based improvisation engine ("improvise like a machine"). The algorithmic responses are based on the analyzed input as well as on internalized knowledge of contextually relevant material. The algorithm fragments MIDI and audio input to short phrases. It then attempts to find a "fit" response by evolving a pre-stored human-generated population of phrases using a variety of mutation and crossover functions over a variable number of generations. At each generation, the evolved phrases are

evaluated by a fitness function that measures similarity to the input phrase, and the least fit phrases in the database are replaced by members of the next generation. A unique aspect in this design is the reliance on a pre-recorded human-generated phrase set that evolves over a limited number of generations. This allows musical elements from the original phrases to mix with elements of the real-time input to create hybrid, and at times unpredictable, responses for each given input melody. By running the algorithm in real-time, the responses are generated in a musically appropriate timeframe.

Approximately forty melodic excerpts of variable lengths and styles were used as an initial population for the genetic algorithm (GA). The phrases were recorded by a jazz pianist who improvised in a similar musical context to that in which the robot was planned to perform. Having a distinctly “human” flavour, these phrases provided the GA with a rich initial pool of rhythmic and melodic “genes” from which to build its own melodies. This is notably different from traditional genetic algorithmic approaches in computer music, in which the starting population is generated stochastically. A similarity measure between the observed input and the melodic content of each generation of the GA was used as a fitness function. The goal was not to converge to an “ideal” response by maximizing the fitness metric (which could have led to an exact imitation of the input melody), rather to use it as a guide for the algorithmically generated melodies. By varying the number of generations and the type and frequency of mutations, certain characteristics of both the observed melody and some aspects of the base population could be preserved in the output.

Dynamic Time Warping (DTW) was used to calculate the similarity measure between the observed and generated melodies. A well-known technique originally used in speech recognition applications, DTW provides a method for analyzing similarity, either through time shifting or stretching, of two given segments whose internal timing may vary. While its use in pattern recognition and classification has largely been supplanted by newer techniques such as Hidden Markov Models, DTW was particularly well suited to the needs of this project, specifically the task of comparing two given melodies of potentially unequal lengths without referencing an underlying model. We used a method similar to the one proposed by (Smith, McNab et al. 1998), deviating from the time-frame-based model to represent melodies as a sequence of feature vectors, each corresponding to a note. Our dissimilarity measure assigns a cost to deletion and insertion of notes, as well as to the local distance between the features of corresponding pairs. The smallest distance over all possible temporal alignments was chosen, and the inverse (the “similarity” of the melodies) was used as the fitness value. The local distances were computed using a weighted sum of four differences: absolute pitch, pitch class, log-duration, and melodic attraction. The individual weights are configurable, each with a distinctive effect upon the musical quality of the output. For example, higher weights on the log-duration difference led to more precise rhythmic matching, while the pitch-based differences led to outputs that more closely mirrored the melodic contour of the input. Melodic attraction between pitches was calculated based on the Generative Theory of Tonal Music model (Lerdahl and Jackendoff 1983). The relative balance between the local distances and the temporal deviation cost had a pronounced effect upon the output – a lower cost for note insertion/deletion led to a highly variant output. Through substantial experimentation, we arrived at a handful of effective configurations.

The computational demands of a real-time context required significant optimization of the DTW, despite the relatively small length of the melodies (typically between two and thirty

and transposition operations. In the end, we had seven mutation functions and two crossover functions available, any combination of which was allowed through a configurable interface.

In order for Haile to improvise in a live setting, we developed a number of human-machine interaction schemes driven by capturing, analyzing, transforming, and generating musical material in real-time. Much like a human musician, Haile was programmed to “decide” when to play, for how long, when to stop, and what notes to play in any given musical context. Our goal was to expand on the simple call-and-response format, creating autonomous behaviour in which the robot interacts with humans by responding, interrupting, ignoring, or introducing new material that aims to be surprising and inspiring. The system received and analyzed both MIDI and audio information. Input from a digital piano was collected using MIDI while the MSP object `pitch~` was used for pitch detection of melodic audio from acoustic instruments. In an effort to establish Haile’s listening abilities in live performance settings, simple interaction schemes were developed that do not use the genetic algorithm. One such scheme was direct repetition of human input, in which Haile duplicated any note that was received from MIDI input. In another interaction scheme, the robot recorded and played back complete phrases of musical material. A simple chord sequence caused Haile to start listening to the human performer, and a repetition of that chord caused it to play back the recorded melody. Rather than repeating the melody exactly as played, Haile utilized a mechanism that stochastically added notes to the melody, similarly to the density mutation function described above.

The main interaction scheme used with the genetic algorithm was an adaptive call-and-response mechanism. The mean and variance of the inter-onset times in the input was used to calculate an appropriate delay time; then if no input was detected over this period, Haile generated and played a response phrase. In other interaction schemes, developed in an effort to enrich the simple call-and-response interaction, Haile was programmed to introduce musical material from a database of previous genetically modified phrases, interrupt human musicians with responses while they are playing, ignore human input, and imitate melodies to create canons. In the initial phase of the project, a human operator was responsible for some real-time playback decisions such as determining the interaction scheme used by Haile. In addition, the human operator of the system triggered events, choosing among the available playback modes, decided between MIDI and audio input at any given time, and selected the different types of mutation functions for the genetic algorithm. In order to facilitate a more autonomous interaction, an algorithm was then developed that chooses between these higher-level playback decisions based on the evolving context of the music, thus allowing Haile to react to musicians in a performance setting without the need for human control. Haile’s autonomous module involved switching between four different playback modes: call-and-response (described above), independent playback, canon mode, and solo mode. During independent playback mode, Haile introduced a previously generated melody from the genetic algorithm after waiting a certain period of time. Canon mode employed a similar delay, but here Haile repeated the input from a human musician. If no input was detected for a certain length of time, Haile entered solo mode, where it continued to play genetically generated melodies until a human player interrupted the robotic solo. Independently of its playback mode, Haile decided between inputs (MIDI or audio) and changed the various parameters of the genetic algorithm (mutation and crossover types, number of generations, amount of mutation, etc.) over time. The human performers did not know who Haile was listening to or exactly how Haile will

respond. We feel this represents a workable model of the structure of interactions that can be seen in human-to-human musical improvisation.

Two compositions were written for the system and performed in two separate concerts. In the first piece, titled "Svobod," a piano and a saxophone player freely improvised with a semi-autonomous robot (see video excerpts - <http://www.coa.gatech.edu/~gil/Svobod.mov>). The second piece, titled "iltur for Haile," involved a more defined and tonal musical structure utilizing genetically driven and non-genetically driven interaction schemes, as the robot performed autonomously with a jazz quartet see video excerpts - <http://www.coa.gatech.edu/~gil/iltur4haile.mov>).

Figure 14. Human players interact with Haile as it improvises based on input from saxophone and piano in "iltur for Haile"

6. Preliminary Evaluation

In order to evaluate our approaches in design, mechanics, perception, and interaction we conducted a user study where subjects were asked to interact with Haile, to participate in a perceptual experiment, and to fill a questionnaire regarding their experience. The study addressed only the first phase of the project, and included only rhythmic applications (user studies for the second phase of the project will be conducted in the future). The 14 undergraduate students who participated in the study were enrolled in the percussive ensemble class at Georgia Tech in Spring 2006 and had at least 8 years of experience each in playing percussion instruments. This level of experience was required to support the musical interaction with Haile as well as to support a meaningful discussion about subjects' experience. Each subject spent about 20 minutes experimenting with four different

interaction modes – imitation, stochastic transformation, perceptual accompaniment, and perceptual transformation. Subjects were then asked to compare their notion of rhythmic stability with Haile's algorithmic implementation. As part of the perceptual experiment on stability, subjects were asked to improvise a one-measure rhythmic phrase while Haile provided a 4/4 beat at 90 BPM. Subjects were then randomly presented with three transformations of their phrase: a less stable version, a version with similar stability, and a more stable version. The transformed measures were generated by the Max/MSP stability external (see section 5.1) using stability ratings of 0.1, 0.5, and 0.9 for less, similar, and more stability, respectively. All phrases, including the original, were played twice. Students were then asked to indicate which phrase, in their opinion, was less stable, similar, or more stable in comparison to the original input. Stability was explained as representing the "predictability of" or "ease of tapping one's foot along with" a particular rhythm. The goal of this experiment was not to reach a definite well-controlled conclusion regarding the rhythmic stability model we used, but rather to obtain a preliminary notion about the correlation between our algorithmic implementation and a number of human subjects' perception in an interactive setting. The next section of the user study involved a written survey where subjects were asked to answer questions describing their impression of Haile's physical design, mechanical operation, the different perceptual and interaction modes, as well as a number of general questions about human-robot interaction and "robotic musicianship". The survey included 39 questions such as: "What aspects of the design and mechanical operation make Haile compelling to play with?" "What design aspects are problematic and require improvements?" "What musical aspects were captured by Haile in a satisfactory manner?" "What aspects were not captured well?" "Did Haile's response make Replica with Musical sense?" "Did the responses encourage you to play differently than usual and in what ways?" "Did the interaction with Haile encourage you to come up with new musical ideas?" "Do you think that new musical experiences, and new music, can evolve from musical human-robot interaction?"

Most subjects addressed Haile's physical design in positive terms, using descriptors such as "unique", "artistic", "stylized", "organic", and "functional". Other opinions included "the design offered a feeling of comfort", "the design was pleasing and inviting, and "if Haile was not anthropomorphic it would not have been as encouraging to play with". When asked about caveats in the design several subjects mentioned "too many visible electronics" "exposed cabling" and suggested that future designs should be "less cluttered." Another critique was that "the design did not appear to be versatile for use with other varieties of drums." Regarding Haile's mechanical operation, subjects provided positive comments regarding the steadiness and accuracy of the left hand and the speed and "smoothness" of the right hand. The main mechanical caveats mentioned were Haile's limited timbre and volume control as well as the lack of larger and more visual movements. Only one respondent complained about the mechanical noise Haile produces. In the perceptual rhythmic stability study, half of the respondents (7/14) correctly identified the three transformations (in comparison, a random response would choose 2.3/14 correctly on average). The majority of confusions were between similar and more stable transformations and between similar and less stable transformations. Only 3 responses out of the total 42 decisions confused a more stable version for a less stable version, implying that larger differences in algorithmic stability ratings made differentiation easier. Only one subject labelled all three generated rhythms incorrectly. Subjects' response to the four interaction

modes was varied. In Imitation Mode respondents mentioned Haile's "accuracy," and "good timing and speed" as positive traits and its lack of volume control as a caveat. Responses to the question "How well did Haile imitate your playing?" ranged from "pretty well" to "amazingly well." Some differences between the interaction modes became apparent. For example, in Stochastic Transformation Mode (STM), about 85% of the subjects provided a clear positive response to the question "Was Haile responsive to your playing?" Only about 40% gave such a clear positive response to this question in Perceptual Accompaniment Mode (PAM). Respondents refer to the delay between user input and robotic response in PAM as the main cause for the "less responsive feel." To the question "Did Haile's responses encourage you to play differently than usual?" 50% of the subjects provided a positive response in STM while only 30% gave a positive response to this question in PAM. When asked to describe how different than usual their playing was in STM, subjects focused on two contradicting motivations: Some mentioned that they played simpler rhythms than usual so Haile could transform them easily and in an identifiable manner. Others made an effort to play complex rhythms to challenge and test Haile's abilities. These behaviours were less apparent in PAM. While only 40% (across all interaction modes) provided a positive answer to the question "Did Haile's responses encourage you to come up with new musical ideas?", more than 90% percent of participants answered positively to the question "Do you think that new musical experiences and new music, can evolve from human-machine musical interactions?", strengthening their answers with terms such as "definitely," "certainly," and "without a question".

Based on the experiment and survey, we feel that our preliminary attempt at Robotic Musicianship provided promising results. The most encouraging survey outcome, in our opinion, was that subjects felt that the human-machine collaboration established with Haile did, on occasions, lead to novel musical experiences and new musical ideas that would not have been conceived by other means. It is clear, though, that further work in mechanics, perception, and interaction design is required to create a robot that can truly demonstrate "musicianship." Nearly all subjects addressed Haile's design in positive terms, strengthening our assumption that the wood and the organic look function well in a drum circle context. Our decision to complement the organic look with exposed electronics was criticized by some subjects, although we feel that this hybrid design conveys the robotic functionalities and reflects the electroacoustic nature of the project. Mechanically, most subjects were impressed with the speed and smoothness of Haile's operation. Only one subject complained about the noise produced by the robot, which suggests that most players were able to either mask the noise out or to accept it as an inherent and acceptable aspect of human-robot interaction. Several subjects, however, indicated that Haile's motion did not provide satisfactory visual cues and could not produce adequate variety of loudness and timbre. The control mechanism for the left arm was developed subsequently to the user study and is currently providing a wider dynamic range. The user study and survey also provided encouraging results in regards to Haile's perceptual and interaction modules. The high percentage of positive responses about the Imitation Mode indicates that our low-level onset and pitch detection algorithms were effective. In general, a large majority of the respondents indicated that Haile was responsive to their playing. Perceptual Accompaniment Mode (PAM), however, was an exception to this rule, as subjects felt Haile was not responding to their actions with acceptable timing. PAM was also unique in the high percentage of subjects who reported that they did not play differently in comparison to

playing with humans. We explain this results by the synchronous accompaniment nature of PAM, which is more familiar to most percussion students. Most subjects, on the other hand, felt compelled to play differently than usual in sequential call-and-response modes such as Stochastic Transformation Mode (STM). Here subjects changed their usual drumming behaviours either by simplifying their rhythms to better follow Haile's responses or by playing complex rhythms in an effort to challenge the robot's perceptual and mechanical abilities. We believe that these behaviours were caused by the novelty effect as players attempted to explore Haile's physical and cognitive boundaries. We assume that subjects may develop more complex interaction behaviours if given longer play times. Given the high level of variance in the notion of rhythmic stability in human perception we feel that our rhythm stability experiment performed better than expected. Some caveats in our method may have also hindered the results. For example, misalignment of subject drumming with the metronome during recording led to misaligned transformations, which may have been unjustifiably perceived as unstable. Also, since the transformed rhythms were generated based on subjects' input, the relative difference between the output stabilities in some cases became minimal and difficult to identify. For example, when a subject's original phrase was extremely stable the algorithm would not be able to produce an identifiably "more stable" phrase. Asking subjects to play a unified mid-stability rhythm as input could have solved this problem, although we were specifically interested in evaluating Haile's perception in a live improvisatory context. As indicated above, the most encouraging results were that 40% of subjects stated that the interaction with Haile encouraged them to come up with new musical ideas and more than 90% claimed that they believe that new musical experiences, and new music, can evolve from such human-machine interaction. This may indicate that although the potential for creating novel musical experiences between humans and robots was not fully realized in our current implementation, the experience led a large majority of the subjects to feel optimistic about the prospect of achieving such novel musical experiences in the future.

7. Future Work

The preliminary user feedback stressed the importance of large visual motions for enabling humans to synchronize with and anticipate the robot's actions. In an effort to address this need, particularly for pitched instruments, we plan to develop a new robotic marimba player that will use several mallets with large and visible striking motions, both horizontally and vertically. Inspired by common human playing techniques, the robot will consist of four arms, each with three degrees of freedom, and a span of one octave (see Figure 15). The robotic arms will be 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. Due to its four independent arms, the robot will only limited by the speed at which each mallet can move, although a certain amount of coordination between neighbouring arms will be required to avoid collisions in the shared workspace. The independent operation of each arm will enable the robot to play sophisticated note combinations faster and more accurately than humans. In an effort to extend the perception, improvisation and interaction capabilities of the robot we plan to develop new algorithmic models for melodic tension, attraction, and

similarity as well as new models for sequential and synchronous musical human-robot interaction. We are also working on a number of interactive improvisational algorithms based on fractals and cellular automata and intend to conduct user studies that will lead to workshops and concerts with the new robot.

Figure 15. The robotic marimba player will have of four arms, each with three degrees of freedom. A servomotor mounted at the end of each arm will rotate the mallets in a vertical plane to strike the bars. An additional servomotor and a linear slide assembly will act together to position the mallets over an octave range

Figure 16. The pivot point of each mallet will be able to move 14 cm towards and away from the marimba to play upper and lower keys, and the mallets will rotate laterally by 90 degrees to reach a one-octave span

8. Acknowledgments

I would like to thank Scott Driscoll, who took a significant role in designing and developing Haile and its rhythmical applications in Phase 1, as well Mark Godfrey, Alex Rae, and John Rhoads, who helped in developing the melodic application in Phase 2. I would also like to acknowledge Georgia Tech's College of Architecture, Music Department and the GVV centre for their support.

9. References

- Baginsky, N. A. (Accessed 2007). The Three Sirens: A Self Learning Robotic Rock Band. <http://www.the-three-sirens.info/binfo.html>
- Biles, J. A. (1999). Life with GenJam: interacting with a musical IGA. *Processions of IEEE Systems, Man, and Cybernetics Conference*, pp. 652-656, ISBN: 0-7803-5731-0, Tokyo, Japan, October 1999
- Brooks, A. G.; Gray J.; Hoffman, G.; Lockerd, A.; Lee, H.; & Breazeal C. (2004). Robot's play: interactive games with sociable machines. *Computer Entertainment*, Vol. 2, No. 3, 10-20
- Chida, K.; Okuma, I.; Isoda, S.; Saisu, Y.; Wakamatsu, K.; Nishikawa, K.; Solis, J.; Takanobu, H.; Takanishi, A. (2004). Development of a new Anthropomorphic Flutist Robot WF-4. *Proceedings of IEEE International Conference on Robotics and Automation*. pp. 152-157, ISBN: 0-7803-8232-3, Barcelona, Spain, April 2004
- Cope, D. (1996). *Experiments in Music Intelligence*, A-R Editions, ISBN 0-89579-256-7 Madison WI
- Coren, S.; Ward, L.M.; & Enns. G. 2003. Sensation and Perception. John Wiley & Sons Inc, ISBN: 0-4712-7255-8, Hoboken, New Jersey
- Cycling74 (Accessed 2007). Max/MSP Web site, <http://www.cycling74.com/products/maxms>.
- Dannenbergh, R. (1984). An On-line Algorithm for Real-Time Accompaniment. *Proceedings of International Computer Music Conference*, pp. 193-198, Paris, France
- Davies, M.E.P.; Plumbley, M.D. Beat tracking with a two state model. *Proceedings of the 2005 IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 241-244, Philadelphia, Penn., USA, 2005.
- Desain, P.; & H. J. Honing (2002). Rhythmic stability as explanation of category size. *Proceedings of the International Conference on Music Perception and Cognition* Sydney, CD-Rom, Australia, July 2007
- Dorssen, M. V. (Accessed 2007). The Cell website. <http://www.cell.org.au/>
- Foote, J. & S. Uchihashi (2001). The Beat Spectrum: a new approach to rhythmic analysis. *Proceedings of IEEE International Conference on Multimedia and Expo*, pp. 881-884, August 2001, Tokyo, Japan
- Johnson-Laird, P. N. (2002). How Jazz Musicians Improvise. *Music Perception*, Vol. 19, No. 3, Spring 2002 415-442
- Jordà, S. (2002). Afasia: The Ultimate Homeric One-man multimedia- band. *Proceedings of the International Conference on New Interfaces for Musical Expression*, pp.1-6, Dublin, Ireland, May 2002
- Kapur, A. (2005). A History of Robotic Musical Instruments. *Proceedings of the International Computer Music Conference*, pp. 21-28, Barcelona, Spain, September 2005

- Lerdahl, F. and R. Jackendoff (1983). *A Generative Theory of Tonal Music*. MIT Press. ISBN 0-2626-2107-X Cambridge, MA
- Lewis, G. (2000). Too Many Notes: Computers, Complexity and Culture in Voyager. *Leonardo Music Journal*, Vol. 10, 33-39
- MakingThings (Accessed 2007). Teleo web site, <http://www.makingthings.com/teleo/>
- Moroni, A. (2000). Vox Populi: An Interactive Evolutionary System for Algorithmic Music Composition. *Leonardo Music Journal* Vol. 10, 49-54.
- Orio, N.; S. Lemouton & Schwarz. D. (2003). Score Following: State of the Art and New Developments. *Proceedings of International Conference on New Interfaces for Musical Expression*, pp. 36-41, Montreal, Canada, May 2003
- Pachet, F. (2003). The continuator: Musical interaction with style. *Journal of New Music Research*, Volume 32, Issue 3, 333 - 341
- Pressing, J. (1994). *Compositions for Improvisers: An Australian Perspective*, Trobe University Press, ISBN: 1-8632-4415-8, Melbourne, Australia
- Puckette, M.; Apel, T. & Zicarelli, D (1998). Real-time Audio Analysis Tools for Pd and MSP. *Proceedings of the International Computer Music Conference*, pp. 109-112, Ann Arbor, Michigan, October 1998
- Rae, G. W. (2005). Robotic Instruments web site
http://logosfoundation.org/instrum_gwr/automatons.html
- Rowe, R. (1992). *Interactive Music Systems: Machine Listening and Composing*. MIT Press, ISBN: 0262181495, Cambridge, MA
- Scheirer, E. (1998). Tempo and beat analysis of acoustic musical signals, *Journal of the Acoustical Society of America*, Vol. 103, No. 1, January 1998, 588-601
- Singer, E.; Feddersen J.; Redmon, C. & Bowen B. (2004). LEMUR's Musical Robots. *Proceedings of International Conference on New Interfaces for Musical Expression*, pp. 181-184, Hamamatsu, Japan, June 2004
- Smith, L., R. McNab, et al. (1998). Sequence-based melodic comparison: A dynamic-programming approach. *Computing in Musicology*, Vol. 11, 101-128
- Tanguiane, A. (1993). *Artificial Perception and Music Recognition*. Springer-Verlag, ISBN: 3-5405-7394-1, Berlin, Germany
- Tokui, N. & H. Iba (2000). Music Composition with Interactive Evolutionary Computation. *Proceedings of the International Conference on Generative Art*, CD-Rom, Milan, Ital
- Toyota (2004). The Trumpet Robot website, <http://www.toyota.co.jp/en/special/robot>
- Trimpin (2000). *SoundSculptures Exhibition: Five Examples*. Munich, Germany, MGM MediaGruppe Munchen.
- Vercoe, B. (1984). The Synthetic Performer in the Context of Live Performance, *Proceedings of the International Computer Music Conference*, pp. 199-200, Paris, France.
- Weinberg, G. (2005). Interconnected Musical Networks - Toward a Theoretical Framework. *Computer Music Journal*, Vol. 29, No. 2, pp. 23-39
- Winkler, T. (2001). *Composing Interactive Music: Techniques and Ideas Using Max*, MIT Press, ISBN: 0262731398, Cambridge, MA

Human Robot Interaction

Edited by Nilanjan Sarkar

ISBN 978-3-902613-13-4

Hard cover, 522 pages

Publisher I-Tech Education and Publishing

Published online 01, September, 2007

Published in print edition September, 2007

Human-robot interaction research is diverse and covers a wide range of topics. All aspects of human factors and robotics are within the purview of HRI research so far as they provide insight into how to improve our understanding in developing effective tools, protocols, and systems to enhance HRI. For example, a significant research effort is being devoted to designing human-robot interface that makes it easier for the people to interact with robots. HRI is an extremely active research field where new and important work is being published at a fast pace. It is neither possible nor is it our intention to cover every important work in this important research field in one volume. However, we believe that HRI as a research field has matured enough to merit a compilation of the outstanding work in the field in the form of a book. This book, which presents outstanding work from the leading HRI researchers covering a wide spectrum of topics, is an effort to capture and present some of the important contributions in HRI in one volume. We hope that this book will benefit both experts and novice and provide a thorough understanding of the exciting field of HRI.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Gil Weinberg (2007). Robotic Musicianship - Musical Interactions Between Humans and Machines, Human Robot Interaction, Nilanjan Sarkar (Ed.), ISBN: 978-3-902613-13-4, InTech, Available from:
http://www.intechopen.com/books/human_robot_interaction/robotic_musicianship_-_musical_interactions_between_humans_and_machines

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2007 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen