

Musicianship for Robots with Style

Marcelo Gimenes

marcelo.gimenes@plymouth.ac.uk

Eduardo Reck Miranda

eduardo.miranda@plymouth.ac.uk

Chris Johnson

c.johnson@plymouth.ac.uk

Interdisciplinary Centre for Computer Music Research,
School of Computing, Communications and Electronics
University of Plymouth, UK
Tel: +44 (0)1752 232579

ABSTRACT

In this paper we introduce a System conceived to serve as the "musical brain" of autonomous musical robots or agent-based software simulations of robotic systems. Our research goal is to provide robots with the ability to integrate with the musical culture of their surroundings. In a multi-agent configuration, the System can simulate an environment in which autonomous agents interact with each other as well as with external agents (e.g., robots, human beings or other systems). The main outcome of these interactions is the transformation and development of their musical styles as well as the musical style of the environment in which they live.

Keywords

Musicianship, artificial life, musical style.

1. INTRODUCTION

Imagine somewhere in the future you get home late at night and begin a conversation with your recently acquired most up-to-date home appliance, your new robot musician pal:

- Hi, Bot, how have you been today?

"Bot" belongs to a new series of robots specially designed to creatively interact with humans and comes standard with a pianist player function. You could have it playing electric guitar as well! Bot was just anxious to report its' new discoveries:

- Not bad at all. I listened to some music on the Internet and practiced my performing skills. Picked up some of those cool jazz pianists. These guys rock! Funny how their melodic lines look pretty much like the romantic pianists. Oh, I also enjoyed playing with Helena, remember, that pianist we met over the Internet last week? She has been doing a lot of progress on improvising standards. I wish one day I will get to this level of expertise!

The movie industry has already shown us scenes of this genre. Based on Asimov and Silverberg's works ([1], [2]), "The Bicentennial Man" [3], for instance, tells the story of Andrew, an android bought by a wealthy family in the year 2005. At first Andrew is no more than a simple housekeeping robot but, after various upgrades and modifications over the time span of two hundred years, he gradually gets human characteristics until he is

finally declared human "in all respects". Likewise, in "Artificial Intelligence" [19], David is a child android that has feelings and, as Pinocchio did, desperately seeks to become a real boy.

These are just two examples out of a series of successful movies that have inspired the general public's dream of upcoming robots that would not only possess external human characteristics but also highly specialized mental abilities and emotions, even to the point of becoming real mates.

How far are we from building machines that embody intelligence, personality, emotion, desire and, who knows, self-awareness? This is a hard question to answer. Many variables come into play but much progress has already been made in many of the relevant and associated fields such as mechanics and artificial intelligence.

In music, for instance, especially with robotic instruments, from the first piano players of the late 18th century to the current Yamaha Disklavier [23] or the QRS Pianomation [15], and the myriad of new instruments and/or interfaces ([8], [17]), there has been a huge improvement. The same applies to the area of robot performers ([20], [21]). WABOT-2 [9] is usually mentioned as one of the first robot-performers. Twenty years ago it was already able to communicate with a person, read a normal musical score, play an electronic organ and even accompany while listening to someone singing.

However, this is neither the subject of this paper nor of the system that we are introducing here. Robots, besides being able to perform unthinkable tasks compared to the abilities of human beings, can also be placed under the scope of systems that demonstrate 'musicianship': their analysis, performance and composition could be intimately linked to "music theory, computer music, cognition, artificial intelligence and human-computer interaction" [16]. The notion of furnishing machines with musicianship has been brilliantly discussed by Rowe in his book *Machine Musicianship* [16].

Therefore, besides the intricate question of self-awareness ("I read" ... "I also enjoyed" ... "I wish") Bots' statement (above) is rich in other equally challenging issues. In that hypothetical day, it would have been able to perceive in order to interact with music and interaction would have led to the transformation of its internal state. In addition, it would have been able to possess its own musical world view and have the means to compare this view with other peoples' (or robots') music in order to realize where it would place itself in terms of differing musical styles. Also, it would have been able to attribute some sort of value in order to make statements such as 'these guys rock', what translates to a simple 'I liked them very much'.

However complicated this machine may look, many steps have already been taken towards its achievement. Cope's [4] "Experiments in Music Intelligence" is well known for having obtained satisfactory results in the field of musical style. Pachet's

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME07, June 7-9, 2007, New York, NY

Copyright remains with the author(s).

Continuator [13] is known for being able to engage into musical dialogues with a human performer. Weinberg and Driscoll [22] built a robotic percussionist ('Haile') designed to perform tasks such as listening to live human players, analyzing perceptual aspects of their playing and, using the result of this analysis to play along with a human improviser. The works of Miranda and Tikhanoff [12] and Manzolli [10] are also interesting examples of complex systems that integrate several aspects of machine musicianship.

The system that we introduce in this paper follows this trend of integrating in the same machine, perceptive and cognitive abilities. We envisage robots that integrate with the culture of their surroundings. Even though "created equal" or with similar abilities, because of their differing interactions, these machines end up by holding a great variety of tastes and interests, melting somehow into their environment.

We adopted a multi-agent system approach to modelling the development of robotic musicianship. The system introduced in this paper is the mental engine that can drive the musical development of robots. The main rationale for this approach is that we consider interaction between robots as an important component of our work. Although software agent technology has already been used for such a purpose [14] the novelty relies specially in the way in which:

- music is represented in the robots' memory,
- musical style is defined and evolves, and
- musical style evolution is measured.

All these topics will be addressed in the sections below. Lacking a better name for our system, we refer to it simply as the 'System'. This paper describes a software simulation of the System, where robots are implemented as software agents. For this reason, hereinafter the words "robot" and "agents" will be used interchangeably.

2. THE SYSTEM: AN OVERVIEW.

We have already reported some preliminary results in an initial implementation that considered only monophonic music [7]. Figure 1 shows an excerpt of an original melody created by an agent from this system.



Figure 1. Excerpt of a robot's original melody.

The results were quite compelling and we then decided to move on and experiment with much more complex music. As a result, the System was designed with the following major characteristics (detailed more fully in the following sections):

a) **Ability to deal with polyphonic music.** We are working under the scope of piano-improvised music. Due to the construction of the perceptive algorithms (item 1, section 3 below), the System is better fitted to a genre of performance (music texture) in which the pianist uses the left hand to play a series of chords and the right hand to play a melodic line.

The System addresses this type of complex music but accepts music that could be considered a subset of it, for example, a series of chords, a melody or any combination of the two. Obviously any music (genre, etc.) that fits into these categories would generate an optimal response by the System but we are also experimenting with other polyphonic music that goes beyond these constraints in order to consider other aspects of music making and evaluate future improvements to the System.

b) **Feature extraction.** Agents initially perceive the music data stream according to a number of "filters" that correspond to basic human sensitive information and the musical features (e.g., melodic direction or melodic inter-onset intervals) that one is interested to take into account in a simulation. This perception results in a parallel stream of data that is subsequently used for segmentation, storage (memory) and style definition purposes.

c) **Segmentation.** The System relies on the notion that music is perceived as a succession of smaller structures that have a particular musical meaning. Therefore, any given musical stream must be segmented in order to be understood and processed by the memory. The behaviour of each individual parallel stream of data that was perceived by an agent contributes to the segmentation algorithm that takes into account principles taken from Gestalt psychology (Figure 2).

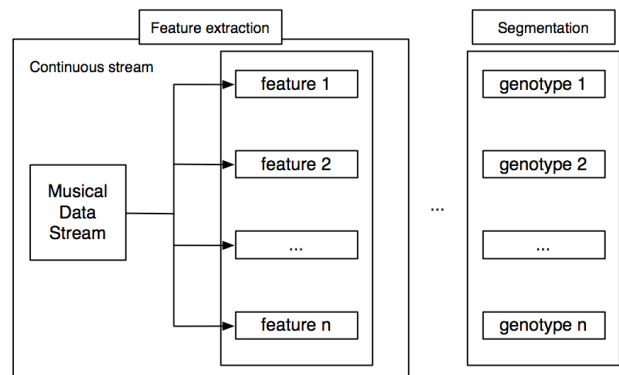


Figure 2. Feature extraction and segmentation.

Each music segment is represented by 'n' "genotypes", in which 'n' is the number of musical features that one wants to configure in a given simulation. By genotype we mean a very basic musical characteristic that is a sub-component of a segment. As an example, the genotype that describes a segment of a melody with 5 notes (4 ascending and one descending) for the musical feature "melody direction" could be: [1, 1, 1, -1]. The initial note is simply a reference and doesn't appear in this genotype. The positive ones are ascending notes and the negative one is the last descending note.

All the genotypes together represent a music segment. The more filter extractors are used the more accurate is the representation. In any segment the number of elements of all the genotypes is the same and corresponds to the number of 'vertical structures' minus one. By "vertical structure" we mean all music elements that happen at the same time.

d) **Memory.** Genotypes are stored in the robot's memory in separate areas ("Feature Tables") according to the category (musical feature) they belong to (section 3.1 below).

e) **General tasks.** Agents perform a number of tasks during their lifetime, the most common being:

1. **Read.** An agent's reading task corresponds to the same ability in human beings. This task is not accomplished in real time, meaning that the agent can use its' ability to read the musical material regardless of the time length of the piece.

2. **Listen.** An agent is able to listen to a musical stream of data that is produced in real time by other agents in the system or by another external musician or system.

3. **Compose.** Agents create new music via a re-combination of the genotypes stored in their memory. This creative process can happen in different modes (section 3.4 below). If the composition is done in "real time", the task is called "improvise". This ability entails the possibility of participation with other agents in the system and/or with an external musician or system.

Every time a musical task is performed the agent's memory is transformed.

f) **Musical style.** We define the musical style of an agent as the state of the agent's memory at any given moment, which comprises all the genotypes, and their related data: dates and number they were read, weights and "connection pointers" (section 3.2 below).

g) **Analysis of style development.** We are experimenting with some measures of distance in order to compare different agent's memories (or the memory of the same agent in different moments) and, therefore, assess the musical style development.

h) **User interface.** The System runs under Mac OS X. Functionalities include opening, saving, importing and exporting several data objects (as well as the whole simulation) such as the agent's memory, music libraries, sets of tasks, compositional and performance maps, etc.

Finally, as a general feature, it is worth mentioning that agents communicate with their environment by receiving and sending streams of data with musical content. We have adopted MIDI in the present version because of its widespread use, as we are interested in promoting integration with other music software. Nevertheless, because the System has its own internal mechanism for symbolically manipulating musical data, we have preserved the possibility of implementing other forms of communication (e.g., Open Sound Control) in the future without compromising any of the System's functionality.

3. THE SYSTEM IN ACTION

The main purpose of the System is to serve as the "musical brain" of robots or software agents. In a multi-agent configuration, the system can simulate an environment in which autonomous agents interact with each other as well as with external agents (e.g., robots, human beings or other systems). The main outcome of these interactions is the transformation and development of these agents' musical styles as well as the musical style of the environment as a whole.

Agents should normally act autonomously and decide if and when to interact. Nevertheless, in the current version of the System we decided to constrain this ability for the sake of better control over the evolution of the musical styles. In this version, agents can choose which music they interact with but not how many times or when they will interact.

In a very basic scenario, a simulation can be designed by simply specifying:

- a) a number of agents,

- b) a number of tasks to each agent and

c) some initial music material for the interactions. Music is placed in the Music Store (Figure 3), a central repository for the compositions.

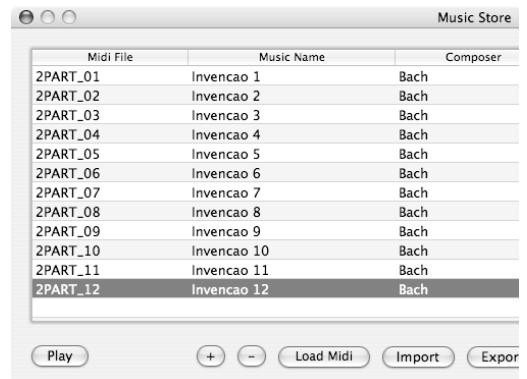


Figure 3. The Music Store.

If one wishes to have more control over the simulation and observe different evolutionary routes, some criteria (such as the name of the composer, year of composition, etc.) can be added to constrain the agents' choices.

If there is some compositional task (compose, improvise and/or collective improvise) at the end of the task the new composition is added to the Music Store.

Once the System is running, it sends the global time step (cycle) to the agents, which then execute the tasks that were scheduled to that time step.

As a general rule, when an agent chooses a piece of music to read (in the form of a MIDI file) or is connected to another agent to listen to its music, it receives a data stream which is initially decomposed into several parallel layers, and then segmented. These steps correspond to the following basic abilities:

1) **Perception.** Agents have specialized "sensors" that perceive the basic characteristics (musical features) of the musical data stream in separate "channels". In other words, the original data stream is decomposed into a number of layers that correspond to the filtered information (see 2.b above).

We have currently implemented 10 filters that follow the melodic line (direction, leap, inter-onset interval, duration and intensity) and non-melodic notes (vertical number of notes, note intervals from the melody, inter-onset interval, duration and intensity).

As it is a difficult task to computationally separate the melodic line (or other voices) from the rest of the notes in a complex musical stream (something that is beyond the goals of the current implementation), we have decided that the melodic line should be defined by the succession of the highest notes above a certain reference note, which is the piano middle C. Anyhow, this can be configured via the user interface.

2) **Segmentation.** As mentioned in section 2.c above, the System relies on the notion that music is perceived as a succession of smaller structures that have a particular musical meaning.

Recently, the term "meme" has been introduced by Dawkins to describe basic units of cultural transmission in the same way that genes, in biology, are units of genetic information. "Examples of memes are tunes, catch-phrases, clothes fashions, ways of making pots or of building arches. Just as genes propagate themselves in

the gene pool by leaping from body to body via sperm and eggs, so memes propagate in the meme pool by leaping from brain to brain via a process which, in the broad sense, can be called imitation." [5]

The idea of employing this expression is attractive because it covers at the same time the concept of structural elements and processes of cultural evolution, something that fits into the purpose of our research. Inspired by Dawkins, we adopted the term "musical meme" or simply "meme" to refer to musical structures in the System. However, note that we do not embrace all aspects of the original memetic theory, as we believe that it needs to be expanded and/or adapted for our purposes.

In the System a Meme is defined as the shortest possible musical structure that cannot be segmented without losing musical meaning. Hence, Memes are generally small structures in the time dimension although they can have any number of notes vertically.

Because the memes were previously separated into channels of basic information, a meme can be described as a series of genotypes, each corresponding to a particular musical feature.

As the description of the segmentation algorithm goes beyond the scope of this paper, we leave it for another opportunity.

3.1 Memory.

According to Snyder [18], "the organisation of memory and the limits of our ability to remember have a profound effect on how we perceive patterns of events and boundaries in time. Memory influences how we decide when groups of events end and other groups of events begin, and how these events are related. It also allows us to comprehend time sequences of events in their totality, and to have expectations about what will happen next. Thus, in music that has communication as its goal, the structure of the music must take into consideration the structure of memory - even if we want to work against that structure".

Informally, the agent's memory comprises a Long Term Memory (LTM) and a Short Term Memory (STM) [18]. The agent's STM is the simplest of the two and stores the 'n' memes ('n' is defined a priori by the user) that were most recently received into the memory. They represent the focus of the "conscious awareness".

A much more complex structure, the LTM is a series of "Feature Tables" (FTs) in which all the genotypes are stored according to their category. FTs are formed by "Feature Lines" (FLs) that keep record of the genotype, the dates of interaction (first reading, last reading and number readings), weight and "connection pointers". In Figure 4 we present the excerpt of a hypothetical FT (for melody leaps) in which there are 11 FLs. The information between brackets in this figure corresponds to the genotype and the numbers in front of the colon correspond to the connection pointers.

Feature n.2 (melody leaps):	
Line 0:	[0 0]: 0 0 0 0 0 0 0 0 0
Line 1:	[2 2 0 1 0 1 2 5 0]: 1
Line 2:	[1 0 0 3 2 2 0]: 2 20 10 10
Line 3:	[1 0 0 0 1 2 2 4]: 3
Line 4:	[2 0 2 0 4 1 3 0]: 4
Line 5:	[0 3 2 7 0 2 0 4]: 5 8 10
Line 6:	[3 0 2 0 3 2 4]: 6 5 3
Line 7:	[1 0 1 2 2 0 3]: 7 3
Line 8:	[2 0 2 0 2 0 0]: 8 31 8
Line 9:	[2 0]: 47 4 9 9 4 9 9
Line 10:	[5 0 8 2 1 2]: 10

Figure 4. A Feature Table excerpt.

3.1.1 Transformation

When a meme is "read" (after perception and segmentation), if the genotype is not present in the corresponding FT, a new FL is created and added to the FT. The same applies to all the FTs in the LTM. The other information in the FLs (dates, weight and pointers) is then calculated.

Figure 5 shows a hypothetical situation in which different genotypes in three different FTs are interconnected via connection pointers.




Figure 5. Interconnection of Feature Tables.

The meaning of these pointers is different depending on the FT. In one of them (chosen by the System's user), called the "first FT", the pointers point to the genotype that was listened sequentially (similar to a Markov chain) in the musical piece. In the other FTs, they point to the index of the genotype in the first FT to which they were connected at the original meme. This information will be used during the execution of the generative tasks (compose and improvise).

The weight of the genotypes increases or decreases depending on whether it is received by the agent's memory during the tasks. The genotype weight is increased every time it is received and decreases if, at the end of the cycle, it is not received. We can say thereafter that the weights represent the relative importance of the genotypes in relation to each other in a FT.

The calculation of the weights is also affected by a "coefficient of permeability" that the agent's memory possesses at any given moment. This coefficient is defined by a group of other variables (attentiveness, character and emotiveness), the motivation being that some tasks entail more or less transformation to the agent's memory depending on the required level of attentiveness (e.g., a reading task requires less attention than an improvisation task). On the other hand, attributes such as character and emotiveness can also influence the level of "permeability" of the memory. These variables are still being tested and deserve a deeper analysis, which will be discussed in a future paper.

3.2 Musical Style

According to Meyer [11] style is "a replication of patterning, whether in human behaviour or in the artefacts produced by human behaviour, that results from a series of choices made within some set of constraints". This patterning can be observed in different scopes, from a single work to a group of works that reveal some common characteristics such as composers, time (period) and/or space (geography).

As seen above, the mechanisms of perception, segmentation and memory transformation embedded in the System constantly affect the state of the memory which is ultimately a complex object comprised by all the genotypes, their weights and connection pointers. The memory can, therefore, be used as the main element to describe the musical style of one (or a group of) piece(s).

3.3 Analysis

This can be accomplished, for instance, in case someone wants to describe the musical style of a particular piece, by designing a simulation in which an agent is forced to read (only once) this piece and nothing more. The musical style could then be extracted from the state of the agent’s memory ultimately revealing the weights and connections between the genotypes that appeared in the piece. The same is applicable to any group of pieces.

The interest of this method may not be immediately visible at first sight but, actually, it turns out to be very useful. For example, we might wish to understand how similar (or different) two compositions are. In this case, the appropriate simulation design would be to have two different agents performing only one reading task each. One agent would read one composition and the other agent would read the other one. At the end of the simulation their memories would then be compared. The mechanism for making such a comparison is discussed in the next section.

From what we have seen up to this point, the state of an agent’s memory describes the accumulated effect of the transformations made on that memory as a result of the execution of a given number of tasks. The state of the memory can therefore be used to describe the agents’ musical worldview.

If we can measure the differences between two memories we can then evaluate how similar or different they are in terms of musical style. At this point wouldn’t we be giving to Bot (remember our robot musician pal from the “Introduction” above?) the necessary tools to help it decide whether those cool jazz pianists really ‘rock’?

An initial attempt to measure the distance between the memories of two agents was implemented as follows. Firstly, for each genotype in turn, the difference in weights of the genotype in the two memories is computed. These differences are then combined using Euclidean distance to produce a measure of distance between the two memories.

Figure 6 shows the evolution of the distance between the musical styles of two agents during 100 cycles in which they listened to the same set of compositions (12 Inventions by J.S. Bach). The agents chose the pieces from this set randomly. During the initial cycles we expected that they could interact with different pieces and that there would be major differences but after a number of cycles the tendency was towards stability.




Figure 6. Style distance between two agents.

Other measures of distance and criteria can be adopted and we are currently experimenting with other possibilities such as:

- the edit (Levenshtein, [6]) distance between the most important (in terms of weight) genotypes and

- the number of connection pointers between the genotypes

3.4 Creative processes

We would like to close this paper with a short description of some of the creative tasks performed by the agents. As mentioned above, the tasks “compose”, “improvise” and “collective improvise” belong to the same genre of creative activities which final product is a composition. The main difference is the fact that “compose” is not performed in real time while the other two tasks are. In addition, their coefficient of permeability can differ based on the fact that a “compose” task could require a smaller degree of attentiveness compared to the other two. This difference, as seen in section 3.1.1 above, can cause an impact in the transformation algorithm.

The general mechanism that generates a new meme works as follows. The agent chooses a genotype from the first Feature Table and then chooses the other genotypes from the other tables in order to re-synthesize a new meme. These are the usual steps:

1. The agent chooses at random the very first genotype of the piece from the first FT. This choice is biased using the distribution of probabilities of all the genotypes in this table.
2. The agent looks at the other FTs and chooses the other genotypes at random but again, based on the distribution of the probability of the genotypes that connect with the genotype chosen in step 1.
3. The agent goes back to the first FT and chooses the next genotype based on the distribution of the probability of happening the “next” genotype in this table (see section 3.1.1 above).
4. The agent repeats step 2.
5. Steps 3 and 4 are repeated as many times as needed until the composition is completed.

The above described “mode of meme generation” is called LTM (for Long Term Memory) generation. Other modes of generation are, nevertheless, designed in the System. For instance, agents can choose memes from the ones that are in their Short Term Memory (STM generation) or from the memes that were re-created beforehand during the composition (meme array generation).

The “mode of meme generation” and other constraints (e.g., note sets, velocity and duration shifts, etc.) are defined in the ‘composition and performance map’, so called because it contains a number of constraints that are valid for the entire composition and performance. As we are dealing with improvised music, composition and performance happen at the same time. Figure 7 shows an excerpt of a ‘compositional and performance map’.

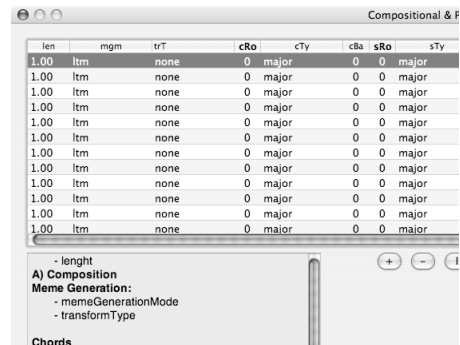


Figure 7. A Compositional and Performance Map excerpt.

4. CONCLUSIONS

We have reported in this paper the implementation of a computer system to control the musical brain of robots and/or software agent-based simulations. Robots and software agents can thereafter integrate with the musical culture of the environment in which they live by transforming and developing their own musical worldview.

Interactions can take place in a purely virtual environment as well as with human beings or other systems. We are testing with different distance measures in order to have an adequate method to assess the robots' musical development.

We are currently looking into the possibility of collaborating with our colleagues of robotic research group to embody the System using humanoid robotic platform that is currently being developed at our University. Future work also includes experimentation with different music styles and human pianists in order to observe the feasibility of the musical interactions in real-life settings.

5. ACKNOWLEDGMENTS

This research is funded by the Brazilian Government's Fundacao Coordenacao de Aperfeicoamento de Pessoal de Nivel Superior (CAPES).

6. REFERENCES

- [1] Asimov, I. and Silverberg, R. *The Positronic Man*. Spectra Ed., 1994.
- [2] Asimov, I. *The Bicentennial Man*. Gollancz, New Ed edition. 2000.
- [3] Columbus, C. (Director). *The Bicentennial Man*. Buena Vista Home Entertainment and Touchstone Studios. 1999.
- [4] Cope, D. *Experiments in Music Intelligence*. A-R Editions, Madison, WI, 1996.
- [5] Dawkins, R. *The Selfish Gene*. Oxford, Oxford University Press, 1989.
- [6] Gilleland, M. Levenshtein distance, in three flavors. Available online at <http://www.merriampark.com/ld.htm> (accessed in January 2007).
- [7] Gimenes, M., Miranda, E. R. and Johnson, C. The development of musical styles in a society of software agents. In *Proceedings of the International Conference on Music Perception and Cognition*, Bologna (Italy), 2006.
- [8] Kapur, A. A History of Robotic Musical Instruments. In *Proceedings of the International Computer Music Conference (ICMC)*, Barcelona (Spain), 2005.
- [9] Kato, I et al. The robot musician "WABOT-2" (WAseda roBOT-2). *ROBOTICS*. Vol. 3, no. 2, pp. 143-155. 1987
- [10] Manzolli, J., Blanchard, J.M. & Verschure, P.F.M.J. (2000) A robot based interactive composition system. In *Proceedings of AAAI*. pp. 435-440. Menlo Park, CA: AAAI Press, 2000.
- [11] Meyer, L.B. (1989). *Style and Music: Theory, History, and Ideology*. Philadelphia: University of Pennsylvania Press.
- [12] Miranda, E. R. and Tikhanoff, V. Musical Composition by Autonomous Robots: A Case Study with AIBO. In *Proceedings of TAROS 2005 (Towards Autonomous Robotic Systems)*. London, UK. (Published by the Imperial College) 2005.
- [13] Pachet, F. The continuator: Musical interaction with style. In *Proceedings International Computer Music Conference*. Goteborg (Sweden), 2002.
- [14] Peters II, R. A., Wilkes, D.M., Gaines, D.M. and Kawamura, K. A Software Agent Based Control System for Human Robot Interaction. In *Proceedings of the Second International Symposium on Humanoid Robots*, 1999.
- [15] QRS. "Pianomation." Available online at <http://www.qrsmusic.com/> (accessed in January 2007)
- [16] Rowe, R. *Machine Musicianship*. MIT Press, Cambridge, MA, 2004
- [17] Singer, E., J. Feddersen, C. Redmon, and B. Bowen. 2004. LEMUR's Musical Robots. In *Proceedings of the 2004 International Conference on New Interfaces for Musical Expression*. Hamamatsu (Japan), 2004.
- [18] Snyder, B. Bob Snyder. *Music and Memory: An Introduction*. Cambridge, Massachusetts: MIT Press. 2000.
- [19] Spielberg, S. (Director). *Artificial Intelligence*. Dreamworks Video. 2001.
- [20] Takanishi, A., Maeda, M., Development of Anthropomorphic Flutist Robot WF-3RIV. In *Proceedings of the International Computer Music Conference*. Michigan, USA, 1998.
- [21] Toyota. "Toyota Partner Robot." Available online at <http://www.toyota.co.jp/en/special/robot/> (accessed in January 2007)
- [22] Weinberg G., Driscoll S. Towards Robotic Musicianship. *Computer Music Journal* 30:4, MIT Press, pp. 28-45. 2006.
- [23] Yamaha, 2007. "Yamaha Disklavier". Available online at <http://www.yamaha.com/> (accessed in January 2007).