TOWARDS AN INTELLIGENT RHYTHMIC GENERATOR BASED ON GIVEN EXAMPLES: A MEMETIC APPROACH

Marcelo Gimenes, *Eduardo Reck Miranda*, *Chris Johnson*[1]
Computer Music Research, School of Computing, Communications and Electronics
University of Plymouth, UK

ABSTRACT

In this paper we present an artificial intelligence system inspired by Richard Dawkin's theory of memes for the composition of rhythmic passages. Rhythmic composition is understood as the process of interconnecting sequences of basic elements (or "rhythmic memes"). The overall design of the system consists of two broad stages: the learning stage and the production stage. In the learning stage, which is the main focus of this paper, the system is trained with examples of musical pieces in order to evolve a "musical worldview". We study the dynamics of this evolution by analyzing the behaviour of the memes logged during the learning and the interaction processes. A series of learning simulations is presented using rhythmic information taken from Brazilian composers of different styles.

Keywords – music representation, rhythm, computer-aided composition, memetics, music cognition

1. INTRODUCTION

From a human perspective, music composition can be regarded as an activity of rearrangement and interconnection of small pieces of musical material or musical structures. Although composers have an almost unlimited range of elements to choose from, they tend to select them, in many cases, influenced by constraints such as cultural background and personal tastes, among others.

Therefore, the organized musical material (to which we can refer as 'pieces of music') follows certain preferences that are the result of these constraints and which are what, after all, defines a 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".

It has been a long-lasting job for musicologists to identify elements that determine the style of composers. If the musical material is organized according to the craftsmanship of the creative mind, on the other hand, it can also be seen as the result of an uncountable number of influences. An interesting musicological analysis is to identify relationships amongst different styles according to the possible influences that composers could have had over each other along their career.

Analysis of musical influences, although inspiring it can be, is a very difficult task because it frequently involves numerous and complex issues. It is true that, many times, structural elements found in the work of different composers can be associated to one another. Most of the times, however, it is hard to state that a specific musical element is inherited from this or that composer.

This network of influences could be placed under a field that we refer to as 'musical ontogenesis'. In Philosophy of Science, ontogenesis refers to the sequence of events involved in the development of an individual organism from its birth to its death. The expression "musical ontogenesis" could thus be used to define the sequence of events involved in the development of the musicality of an individual.

As we said above, it is particularly interesting to study the influence that the musical material perceived by musicians has on the development of their musical knowledge, particularly in composers and improvisers. Intuitively, it should be possible to associate the manner in which composers or improvisers set up musical structures with their educational background. But how objectively can we study this?

Although it may be extremely difficult (in the real world) to establish chains of musical influence in all sorts of music styles, it is an inspiring idea to use mechanisms of musical influence as a tool to generate evolutionary paths of musical creation in artificial worlds. Intelligent agents could learn from each other's music production (as well as from real composers) and evolve their own musical worldviews there from. But what can be evolved and how?

Boroda [1] explains that a 'basic semantic unit' in music is a subject that has been discussed in musicology for more than two centuries. Even though, from the citations below, we can see that there is not much agreement as to what should be considered a basic musical semantic unit:

"The smallest structural unit is the phrase, a kind of musical molecule consisting of a number of integrated musical events, possessing a certain completeness, and well adapted to combination with other similar units. ... The length of a phrase may vary within wide limits. [...] nearly always the phrase crosses the metrical subdivisions, rather than filling the measures completely." ([13] pp. 3-8, cited in [1])

[1] Email: {marcelo.gimenes, eduardo.miranda, c.johnson}@plymouth.ac.uk

"A motif is the smallest part of a musical idea which is both meaningful (expressive) and constructive. ... In the process of motivic continuation parts (particles) within a motif can be sometimes isolated. [...] a phrase is such a unification of motifs which does not built a period. Usually, phrases unite the motifs into pairs". ([10] pp. 552-563, cited in [1])

More recently, Dawkins [3] proposed a different approach to this issue based on Darwin's ideas to explain evolution in cultural domains. According to Dawkins, the 'basic units of cultural transmission' are contained in 'memes' in the same way that genes, in biology, contains units of genetic information. The initial argument of this theory refers explicitly to musical aspects:

"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."([4], p. 206)

As previously mentioned, the definition of the exact length/boundaries of a musical meme is a very complex subject for a number of reasons [7] [8]. Roughly speaking, different individuals can identify different memes in the same or in different pieces of music in accord with, among other factors, their previous personal musical background.

Cox [2] argues that the "memetic hypothesis" is based on the concept that the understanding that someone has on sounds comes from the comparison with the sounds already produced by this person. The process of comparison involves tacit imitation, or memetic participation that is based on the previous personal experience on the production of the sound.

Gabora [5] explains that, in the same way that information patterns evolve through biological processes, mental representations, or memes, evolve through the adaptive exploration and transformation of an informational space through variation, selection and transmission. Our minds perform tasks on its replication through an aptitude landscape that reflects internal movements and a worldview that is continuously being updated through the renovation of memes.

Some rhythmic generating systems have already been proposed [6]. Pachet [12], for instance, describes an evolutionary model where a group of agents play rhythms together in real time without prior knowledge about the music to play. Agents play in cycles to which transformation rules are applied in order to produce new variations.

Our contribution to this field is presented in the next section through RGeme (Rhythmic Meme Generator), a system meant to be used to generate rhythm streams and as a tool to observe how different rhythm styles can originate and evolve in an artificial society of software agents.

2. THE MODEL

RGeme was designed primarily as an artificial intelligence system for the generation of rhythmic streams based on the computational paradigm of intelligent Agents [9]. Agents are able to sense the existence of music compositions in the environment and to choose the ones with which they will interact. Eventually, they will extract the rhythmic information from the chosen compositions and, conversely, generate new rhythmic streams.

At the beginning of a simulation a number of Agents are created. Each agent is associated with an identity (name), a number of tasks ("Goal Matrix") and the criteria ("Evaluation Matrix") that it will apply to choose the compositions with which it will 'interact' ("Candidate Music").

There are three types of agent tasks: listening, practicing and composing. They should be performed during three different stages to which the Agents will go through their lifetime: Listener, Student, Composer. As Listeners, Agents can only execute listening tasks. In the Student stage, Agents can listen to and practice rhythms. In the last stage, as Composers, Agents can execute listening, practicing and composing tasks. Broadly speaking these stages and tasks split the model into two general phases: the learning and the production phases. Evidently, listening and practicing tasks focus mainly on the learning phase whereas the composing task focuses mainly on the production phase.

Before the execution of listening and practicing tasks, the Agents choose the Candidate Music according to the Evaluation Matrix (composer's name and/or year of composition). Agents then parse the Candidate Music in order to extract rhythmic memes (Candidate Memes). In order to keep the model simple, at this stage of our research the rhythmic memes have a fixed length that corresponds to a music bar.

Agents store their musical knowledge in a Style Matrix (SM) in which every entry is related to a unique rhythmic structure (rhythmic meme) with the following information:

- the dates (represented in terms of a counter that calculates each interaction cycle) in which the memes are first and last listened to,
- the number of times the memes are listened to,
- the weight (importance) the memes hold due to the various interactions with the Candidate Memes and
- the Candidate Music the meme is listened from.

Style Matrices also hold 'Composition Maps', which correspond to the ways the Candidate Memes are interconnected in the Candidate Compositions.

Rhythmic memes are being currently coded as vectors whose entries are 0s and 1s (Figure 1), where 1s mean the trigger of sounds and 0s represent rests and time placeholders. One of the drawbacks of this representation is that information such as the position of the meme in the musical stream, the intensity of each sound and the articulation (duration of the sounds) are not taken into account. Nevertheless, this representation is still useful for the initial implementation of the system and can be extended to include the above-mentioned aspects.

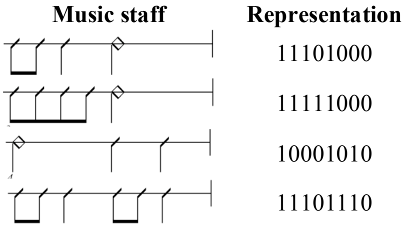

Figure 1. Music staff and corresponding meme representation

Every time a composition is chosen and the Candidate Memes are parsed, a *transformation algorithm* is applied as follows.

- in the beginning of any simulation the SM (Style Matrix) is empty;
- the first parsed Candidate Meme (CM) is copied to the SM (all memes in the SM are called SMMs - Style Matrix memes). The weight of this first SMM is set to 1 and the dates of first and last listening are set according to the general time controlled by the system;
- the next CM is then compared with all the SMMs: (a) if they are different (distance = 1, see below) CM is copied to the SM and its weight is set to 1 and (b) all other SMMs have their weight upgraded according to their distance to CM; and
- the previous step is repeated to all the Candidate Memes.

The distance between two given memes a = [a1, a2, ... an] and b = [b1, b2, ... bn], is defined in equation 1 (block distance):

$$d(a,b) = \frac{1}{n}\sum_{i=1}^{n} | a_i - b_i | \qquad (1)$$

For example,

- the distance between the memes a = 01011101 and b = 11011101 is d(a, b) = 0.125 while
- the distance between the memes a and c = 11111111 is d(a, c) = 0.375.

Once all CM are compared with all SMMs, the new memes are copied and the remaining ones have their weight upgraded, a forgetting effect is applied to the SMM that don't appear in the current Candidate Music.

Up to this point we described how the transformation algorithm alters the musical knowledge of the Agents as a result of the execution of the listening and practicing activities.

Since the learning phase is the main focus of this paper, it suffices to say that in the production phase the Agents execute composition tasks mainly through the reassignment of the various Composition Maps according to the information previously stored in the learning phase. Composition tasks, beyond the production of new material, also have a transformation effect on the Style Matrix where all memes are updated according to the musical material used in the newly produced rhythms.

The model has the potential to execute intricate simulations with several Agents learning at the same time from rhythms by composers from inside and outside the system's environment.

A number of illustrative examples are discussed below.

3. SIMULATIONS

A group of 35 pieces by Brazilian composers Chiquinha Gonzaga, Ernesto Nazareth, Jacob do Bandolim and Tom Jobim was selected. The first three composers are from the end of the 19th and beginning of the 20th century. The latter personifies the 'bossa nova', a rhythm that emerged as a softened version of the 'samba' in the fifties.

We ran a series of simulations from which we selected a few whose summary is shown in Table 1. All agents executed a total of 100 listening tasks and were not allocated any practicing or composing task. Agent A listened to a single piece of music ('Atraente', by Chiquinha Gonzaga) all the time. Agent E listened only to a group of compositions (Jacob do Bandolim and Ernesto Nazareth) during one phase. Agent H listened to a set of compositions by Ernesto Nazareth during the first 50 interactions and to a different set by the same composer during the remainder ones. Agent M listened only to compositions by Jacob do Bandolim during the first 33 tasks, only to compositions by Tom Jobim during the next 33 tasks and again only to compositions by Jacob do Bandolim, during the remainder ones.

Agent	Phases	n. piec.	Composers
A	1	1	Gonzaga
E	1	All	Bandolim and Nazareth
H	2	1 set 1 set	Nazareth, Nazareth
M	3	All All All	Bandolim, Jobim, Bandolim

Table 1. Summary of selected simulations

At each cycle, that corresponds to the execution of one listening task, Agents:

- chose one Candidate Music,
- parsed the rhythmic Candidate Memes and
- applied the transformation algorithm

The system generated a new Style Matrix after the accomplishment of each task. All the resulting Style Matrices were logged in order to observe the behaviour of all SMMs (Style Matrix memes) during the interaction processes.

In Table 2 we show the first of all Style Matrices generated during Agent's E lifetime. It shows how Agent's E musical worldview looked like after it listened to the first music ('Fado Brasileiro', by composer Ernesto Nazareth):

#	SMM	dFL	dLL	nL	W
1	01010110	1	1	2	1.026
2	01011000	1	1	2	1.017
3	11010000	1	1	2	1.021
4	00100010	1	1	2	1.013
5	01110111	1	1	4	1.025
6	11011101	1	1	6	1.022
7	10010111	1	1	6	1.023
8	10010101	1	1	4	1.019
9	11110111	1	1	15	1.014
10	10001000	1	1	1	1.000

SMM: Style Matrix meme
dFL: date of first listening
dLL: date of last listening
nL: number of listening
W: weight

Table 2. Extract from 1st Style Matrix

Notice that each SMM already have different weights that corresponds to the number of times that each one of them were listened to and also to the similarities that they had comparing with the other Candidate Memes. The following Table shows an extract of the next (second) Style Matrix after Agent E listened to the second music ('Apanhei-te Cavaquinho', by composer Ernesto Nazareth):

#	SMM	dFL	dLL	nL	W
1	01010110	1	1	2	1.032
2	01011000	1	1	2	1.016
3	11010000	1	1	2	1.022
4	00100010	1	1	2	1.009
5	01110111	1	1	4	1.039
6	11011101	1	1	6	1.035
7	10010111	1	1	6	1.035
8	10010101	1	1	4	1.025
9	11110111	1	2	19	1.053
10	10001000	1	1	1	0.995
11	00000111	2	2	1	1.021
12	11111111	2	2	28	1.040
13	10000111	2	2	2	1.017
...	...	...	...	...	...

Table 3. Extract from 2nd Style Matrix

Notice, for example, that:

- after the first cycle of interaction (Style Matrix 1 - Table 2), meme 01011000 (second in the list) had been listened (nL) 2 times and its weight (W) was 1.017. After the second cycle of interactions (Style Matrix 2 - Table 3) its number of listening was still 2 but its weight had fallen to 1.016. This was due to the fact that this meme was not present in the second composition and, therefore, Agent E 'slightly forgot' about it.
- meme 9 (11110111), the most prominent in the first music was listened to 15 times in the first interaction. The same meme was listened to only 4 times more in the second music and, for this reason, its weight raised just by a small amount, from 1.014 to 1.053.
- meme 12 (11111111) only appeared in the second composition but because it was listened to 28 times, in the second interaction its weight was already very high, comparing to the others, 1.040.

A set of graphics was generated for each Agent to describe the behaviour of all the SMMs during their lifetime. Obviously, if there is only 1 music to be listened to, as it was the case of Agent A, all the SMMs are learned at once. This was not the case, however, of Agent H, which learned 57 memes up to time (cycle) 72 (Figure 2). Observe that, after a period of stabilization, from time 33 to time 50, a rupture occurs in time 51, when Agent H begins to learn from a different set of compositions.

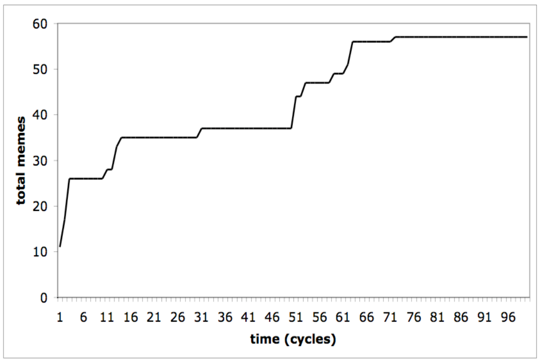

Figure 2. Number of memes learned in time (Agent H)

It was also possible to observe the number of times that each one of the memes was listened to by the Agents at any given time. In Figure 3 we show, for instance, how many times, at the end of the simulation, Agent M had listened to the first learned 10 memes:

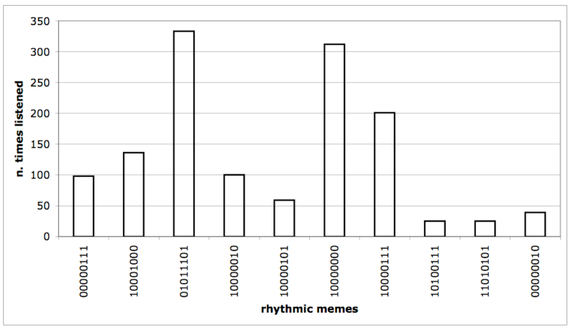

Figure 3. First 10 memes learned by Agent M

Apart from the facts already mentioned above, one of the most interesting features that RGeme is capable to exhibit is the evolution of the importance (weight) that a meme has during the learning phase of an Agent. The increase or decrease of the importance of the memes is the direct result of the number of times and the date they were listened to and/or practiced.

Figure 4 shows the behaviour of all the memes learned during Agent 'M's lifetime. Every time a new meme was learned a new line appeared. If a meme was not heard during a certain time, its curve started to fall ('forgetting effect').

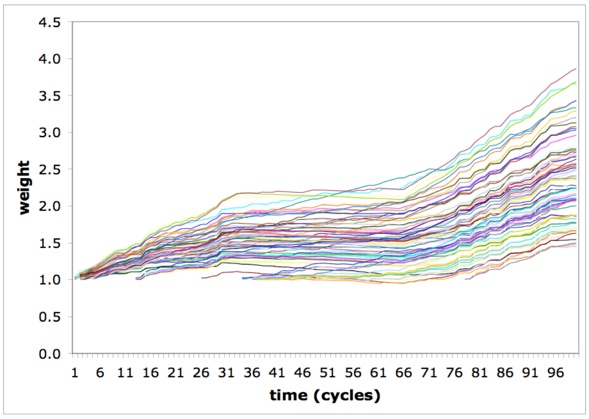

Figure 4. Memes curve of importance in time

As it is obviously very difficult do visualize the evolution of all the 78 memes in the same graph, in Figure 5 we show a selection of a few of them. Figure 6 shows the same memes in a normalized graph after, at any given moment, the weights of all SMM were divided by the highest weight. Some typical behaviour that emerged from the interactions is described in the paragraphs below.

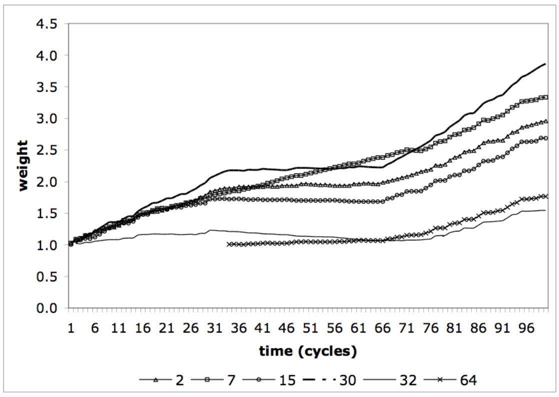

Figure 5. Memes weight in time (selection)

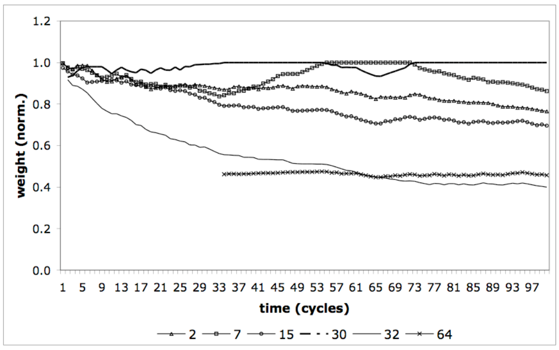

Figure 6. Memes weight in time (normalized)

In Table 4 (ordered by descending weight), each series corresponds to the memes shown in Figures 5 and 6, according to the data collected in the last Style Matrix generated by the system:

Series	SMM	dFL	dLL	nL	W
30	11111111	2	100	513	3.86
7	10000000	1	100	312	3.33
2	00000111	1	100	98	2.96
15	11011000	1	99	112	2.69
64	01011010	34	61	8	1.76
32	00101010	2	95	33	1.54

Table 4. Description of memes

As we can see, meme 30 (11111111) was the winning one. The first time it was heard was in time 2 and, apart from an unstable period at the beginning and at the middle of the simulation, this meme was predominant over the others most of the time.

Meme 7, the second in importance at the end of the simulation, had a relatively stable growth since it was first listened to. In general, the changing of criteria in the Goal Matrix, at cycles 34 and 66, influenced all the memes as we can see by the angle change in the learning curves (Figure 5). We could conclude, therefore, that meme 7, was as important in the music by Jacob do Bandolim as in the music by Tom Jobim.

Memes 2 and 15 had a similar behaviour during Agent's M lifetime but meme 2 won over meme 15, because towards the end of the simulation it was more relevant in the music by Tom Jobim (Table 4).

Finally, just to look at another example from the myriad of different situations that can be observed, although meme 32 begun to be listened to in time 2, its relative importance comparing to the other memes was never very high. On the other hand, meme 64, which appears only in time 34 (music by Tom Jobim), at the end of simulation was victorious over meme 32.

Finally, at the end of the simulation, some of the memes were definitely winners over others: Figure 7 (musical representation) and Table 5 (numeric representation and weight).

Figure 7. Musical representation of the winner memes

#	Meme	Weight
30	11111111	3.86
19	01111111	3.69
4	01011101	3.66
12	11011101	3.65
22	01111101	3.43
54	11111101	3.43

Table 5. Winner memes

As we previously mentioned, all this information will be eventually used by the Agents later on, during their practicing and composing stages, in order to generate new rhythmic material.

4. CONCLUSION

In this paper we presented the learning stages of RGeme, an Agent-based artificial intelligence system inspired by Richard Dawkin's theory of memes for the composition of rhythmic streams.

We presented a series of learning simulations using rhythmic information taken from Brazilian composers of different styles. By analysing the behaviour of the memes logged during the learning and the interaction processes, we studied the dynamics of the evolution of new musical worldviews that will be a major component during the Agent's practicing and composing phase.

We are currently adapting the model to a multi-agent system whereby different Agents will be able to learn from different pieces of music as well as by interacting with each other.

In the future, besides the rhythm information that is being currently employed, the system will also deal with more complex musical structures that consider note information (pitches and vertical structures).

5. ACKNOWLEDGEMENTS

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

6. REFERENCES

[1] Boroda, M. G., "*Towards a Phrase Type Melodic Unit in Music*", Musikometrika 50/4, 15-82, 1992.

[2] Cox, A., "*The mimetic hypothesis and embodied musical meaning*", MusicæScientiæ, V(2):195–212, 2001.

[3] Dawkins, R., "*The Blind Watchmaker*", Penguin Books, London, 1991.

[4] Dawkins, R., "*The Selfish Gene*", Oxford University Press, 1989.

[5] Gabora, L. M., "*The origin and evolution of culture and creativity*", Journal of Memetics – Evolutionary Models of Information Transmission, vol. 1, 1997

[6] Horowitz, D., "*Generating Rhythms with Genetic Algorithms*", Proceedings of the 1994 International Computer Music Conference, San Francisco, ICMA, 1994.

[7] Jan, S., "*Meme Hunting with the Humdrum Toolkit: Principles, Problems and Prospects*", Computer Music Journal (2004) 28(4):68-84.

[8] Jan, S., "*Replicating sonorities: towards a memetics of music*", Journal of Memetics–Evolutionary Models of Information Transmission, vol. 4, 2000.

[9] Maes, P., "*Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*", MIT Press, 1991.

[10] Mazel, L.A., "*Analiz muzykal'nykh proivedenij*" (The analysis of a musical composition), Moskva: Muzyka. 1967.

[11] Meyer, L.B., "*Style and Music: Theory, History, and Ideology*", Philadelphia: University of Pennsylvania Press, 1989.

[12] Pachet, F., "*Rhythms as emerging structures*", Proceedings of 2000 International Computer Music Conference, Berlin, ICMA, 2000.

[13] Schoenberg, A., "*Fundamentals of Musical Composition*", London, Faber and Faber Limited, 1983.