Extracting Patterns from Brazilian Guitar Accompaniment Data

Ernesto Trajano de Lima[1] and **Søren Tjagvad Madsen**[2] and **Márcio Dahia**[3] and **Gerhard Widmer**[4] and **Geber Ramalho**[5]

Abstract. The analysis of expressive performance, an important research topic in Computer Music, is almost exclusively devoted to the study of western classical piano music. Instruments like the acoustic guitar and styles like Bossa Nova and Samba have almost never been studied, despite their harmonic and rhythmic richness. This paper describes some experimental results obtained with the extraction of rhythmic patterns from the guitar accompaniment of Bossa Nova songs. It also discuss the relevance of such experiments, as well as the importance of the computer techniques used in them, to the music and musical culture in general.

1 INTRODUCTION

It is common sense that playing music in the exact way as written in the score results in a mechanical and uninteresting succession of sounds. To make the written music interesting, the musician is required to make variations on various musical parameters, such as: local tempo (*accelerandi*, *ritardandi*, *rubato*); dynamics; notes articulation (*staccati*, ligatures, etc.); micro-silences between the notes, etc. [11].

Studying this phenomenon, also known as *expressive performance*, usually involves discovering some sort of systematic relationship within a piece and/or among pieces. These relationships can be described in many ways, from different points of view or levels of abstraction, and including various musical parameters. Examples of such relationships are rules like *"lengthen a note if it is followed by a longer note and if it is in a metrically weak position"* or *"stress a note by playing it louder if it is preceded by an upward melodic leap larger that a perfect fourth"* [12].

The research on expressive performance today is almost exclusively devoted to the western classical music composed for the piano. We are interested in the study of the *Música Popular Brasileira* (Brazilian Popular Music)—MPB, represented by artists like João Gilberto, Tom Jobim, Caetano Veloso, Gilberto Gil, etc., in particular the guitar accompaniment. There is, however, a fundamental difference between these two genres: While in western classical music there is one "official" notated version of musical pieces (the score), in MPB it does not exist. What is usually available is the chord grid only, and, sometimes, the score for the melody. So, in this genre (MPB), the rhythmic accompaniment must be determined by the performer himself.

It is known that the guitar accompaniment in styles like *Bossa Nova* and *Samba* is built by the concatenation of certain rhythmical patterns [5, 9]. There are, however, several aspects of the accompaniment construction that are only known by practitioners of these styles, i.e., they are still not formally known.

This paper presents an experiment that focus on the discovery of rhythmic patterns in Bossa Nova music. For this, two different performers played several songs on a MIDI guitar, which were processed in the form of strings. Based on a catalogue of Bossa Nova patterns [3]), we tried to identify how much of it could be automatically found in the data we collected[6]. We also discuss the cultural issues behind the experiment.

The remainder of the paper is organized as follows: In Section 2, we describe how the data was acquired and the representation we used. In Section 3, we present the experiment itself, as well as the results we obtained. In Section 4, we discuss the implications and importance such an experiment has for the MPB–related musical culture, and, finally, in Section 5, we present some conclusions and future directions for this work.

2 DATA ACQUISITION AND REPRESENTATION

For the experiment, two different players, referred to as Player1 and Player2, were invited to record the accompaniment of some Bossa Nova songs on a MIDI guitar. Player1 performed the following songs: *Bim Bom*, *O Barquinho*, *Insensatez* (How Insensitive), *Garota de Ipanema* (Girl from Ipanema), *Só Danço Samba*, and *Wave*. From Player2 we recorded *A Felicidade*, *Chega de Saudade*, *Corcovado*, *Desafinado*, *Eu Sei Que Vou Te Amar*, *Samba de uma Nota Só*, *Garota de Ipanema*, *Só Danço Samba*, *Insensatez*, *Tarde em Itapoã*, and *Wave*. In the total, we collected 16 recordings (ca. 30 minutes of music). It was requested the performers to play the songs according to a provided notation (the chord grid as notated by Chediak [2]).

The acquired data was, however, not ready for usage. Probably due to technological restrictions, the resulting MIDI files were noisy, i.e., they contained notes and events that were not played. So, it was necessary a data cleansing step, where these extra events were manually removed. After correction, the data was beat tracked at the eight note level using *BeatRoot* [4], an interactive system that outputs the MIDI file beat tracked.

[1] Centro de Informática (CIn), Universidade Federal de Pernambuco (UFPE), Brazil, email: etl@cin.ufpe.br
[2] Austrian Research Institute for Artificial Intelligence (ÖFAI), and Department of Computational Perception—University of Linz, Austria, email: soren.madsen@ofai.at
[3] CIn, UFPE, Brazil, email: mlmd@cin.ufpe.br
[4] ÖFAI, and Department of Computational Perception—University of Linz, Austria, email: gerhard.widmer@jku.at
[5] CIn, UFPE, Brazil, email: glr@cin.ufpe.br

[6] The catalogue reflects the rhythmic patterns used by João Gilberto, the "inventor" of this style.

As we are interested in the discovery of rhythmic patterns, the exact pitches played (i.e., A, B, C♯, etc.) are not that relevant. A much more relevant abstraction is the right hand finger used by the player to pluck the string[7]. So, we used an algorithm—as described by Trajano et al. [10]—to automatically determine the fingering for each song (Figure 1). Letters T, F, M and R represent, respectively, the thumb, fore, middle and ring fingers, crosses (+) represent the beats, and pipes (|) represent the measure bars. Each beat was equally divided by four, so each letter, cross or minus (–) represents the duration of a 32nd. Except for the last line, that represents exclusively the beats, each of the remaining lines represents one guitar string, ordered from higher to lower (i.e., first line represents the high E string, second line the B string, and so on until the low E string).

```
|----------------|----------------|--
|R---R-----R-----|R---R-----R-----|R-
|M---M-----M-----|M---M-----M-----|M-
|F---F-----F-----|F---F-----F-----|F-    [...]
|----------------|----------------|--
|T-------T-------|T-------T-------|T-
|+---+---+---+---|+---+---+---+---|+-
```

Figure 1. Right hand fingering for song *Insensatez*, as played by Player2

In order the reduce the complexity of the pattern extraction, this initial representation was reduced to a one-dimensional string. This simplified string is formed by the alphabet $\Sigma = \{b, B, p, P, l, a, A, s, S, -, +, |\}$. The meaning of each symbol is the following:

- Uppercase letters stand for events that occur on-beat, while lowercase letters for off-beat events;
- Letter b stands for “bass”, i.e., events played with the thumb only;
- Letter p stands for “chord” (*sic*), i.e., events that are played with some combinations of two or more of fingers F, M and R[8];
- Letter l also stands for “chord”, but a chord whose duration goes beyond the measure it was played (it means that we make a difference between a chord that is completely within a single measure and a chord that starts in one measure and ends in the next one);
- Letter a stands for “all”, i.e., b and p played together;
- Letter s stands for “single note”, i.e., events that are played with only one of fingers F, M and R; and
- Symbols $+$, $-$, and | have the same meaning stated before.

Note that this reduction is also done by the musicians themselves: they usually describe the rhythmic patterns as sequences of “baixos” or *basses* (events played with the thumb only) and “puxadas” or *chords* (events played with some combinations of two or more fingers). Figure 2 depicts part of the fingering for song Insensatez. Above the thick black line is the fingering as output by the right hand fingering algorithm. Under it is the resulting simplified string.

```
|----------------|----------------|--
|R---R-----R-----|R---R-----R-----|R-
|M---M-----M-----|M---M-----M-----|M-
|F---F-----F---F-|F---F-----F-----|F-
|----------------|----------------|--
|T-------T-------|T-------T-------|T-
|+---+---+---+---|+---+---+---+---|+-
=====================================
|A---P---B-p-+-s-|A---P---B-p-+---|A-
```

Figure 2. Fingering and one-dimensional string for song Insensatez, as played by Player1

3 EXPERIMENT

Dahia et al. [3] used a catalogue containing 21 rhythmic patterns of Bossa Nova guitar music, acquired manually from João Gilberto’s performances, to automatically generate the guitar accompaniment of some songs. Figure 3 depicts some of these patterns.

Figure 3. Examples of rhythmic patterns from the catalogue

As one can imagine, the manual acquisition of such a catalogue is tiresome and error prone. It can even be said that, in practice, it can not be done if you want to have the catalogues of many players[9]. So, the acquisition of catalogues of this kind must be done automatically (or semi-automatically, at least). The main objective of our experiment is to try to identify at which extent this Bossa Nova catalogue could be automatically found in the data we collected. This will give us hints on algorithms that could be applied to the acquisition of rhythmic catalogues for other styles, as Samba, for instance.

For this experiment, we used three algorithms: Boyer–Moore [1], FlExPat [8] and SimilaritySegmenter [6][10]. We also rewrote the catalogue in the form of the simplified string previously described.

The first algorithm we used was Boyer–Moore. We tested occurrences of each pattern in the data set, and, from the 21 patterns in the catalogue, the algorithm could only find one pattern (P1) in the whole data set, appearing only in three of the songs. Table 1 summarizes the results. Each entry in the table represents the number of occurrences of P1.

The main question now is the following: is there any flaw in the methodology we used? Since João Gilberto is the inventor of the Bossa Nova guitar style, we expected to find a more representative

[7] We assume the player is right handed.

[8] The terms “baixo” and “puxada” may explain more clearly the origin of letters b and p!

[9] See, for instance, the work of Owens [7]. It took him 16 years to build a similar catalogue for Charlie Parker *only*.

[10] The detailed descriptions of these algorithms is beyond the scope of this paper. Please, refer to the original articles for a detailed description.

Table 1. Number of occurrences of pattern P1

Song	Player1	Player2
Desafinado	0	1
Garota de Ipanema	0	6
Insensatez	4	29

number of the patterns in the data set. Besides, by hearing the songs it is possible to perceive other patterns from the catalogue.

Actually, the main problem is that an exact matcher, such as the Boyer–Moore algorithm, isn't suitable for the task, since it is not able to cope with any sort of deviations (small anticipations and delays of events, for instance) that are usually done by the performers. Figure 4 depicts some of these variations (the "s" above the last 16th represents a single note).

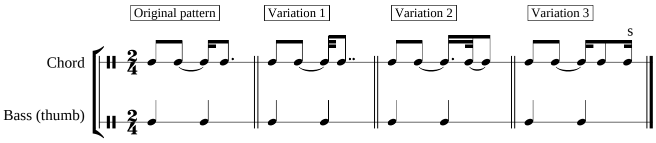

Figure 4. Some variations of part of a pattern

So, in order to discover patterns, we must use algorithms that can perform the so called *inexact string matching*. For that, we used FlExPat and SimilaritySegmenter[11].

In the case of FlExPat, the configuration we used was the following:

- $m_{min} = 17$ (minimal pattern length);
- $m_{max} = 34$ (maximal pattern length);
- similarity threshold: 0.75 (normalized values);

Lengths m_{min} and m_{max} correspond to patterns varying from one to two measures. It does not mean, however, that the patterns necessarily start at the beginning of the measure. There is no way to specify such a constraint in FlExPat. On average, each song has 48 classes of patterns (smallest number of classes was 40 and greatest number was 78), which is an acceptable number.

Table 2 summarizes the results we obtained for FlExPat.

Although the algorithm found several interesting patterns in the data set, there were some problems with the results we obtained. The main problem we found was that the extracted patterns, many times, start from the middle of a measure, such as in `--B---B---+---|P---B---B---+---|P` (FlExPat also found the expected pattern, `|P---B---B---+---|P---B---B---+--`). This brings some problems to the evaluation of the results: due to the great number of patterns with this structural malformation, it becomes difficult to validate the patterns. This problem is even bigger when most of a class is formed by such patterns (it happened many times, unfortunately).

[11] Note that now, with the inexact matching, there is a difference in our methodology: while with the Boyer–Moore algorithm we tested the existence of the patterns of the catalogue directly in the data set, now we must first induce patterns and only then look and see if the induced patterns have some resemblance with the patterns in the catalogue.

Table 2. Occurrences of patterns in data set (FlExPat).

Pattern	Player1	Player2	Pattern	Player1	Player2
P1	*yes*	*yes*	P12	no	no
P2	*yes*	no	P13	no	no
P3	no	no	P14	no	no
P4	*yes*	no	P15	*yes*	no
P5	no	no	P16	no	no
P6	no	no	P17	no	no
P7	no	no	P18	no	no
P8	no	no	P19	no	no
P9	*yes*	*yes*	P20	no	no
P10	*yes*	*yes*	P21	no	no
P11	no	no	/	/	/

For SimilaritySegmenter, we allowed a maximal number of 2 mismatches per measure. We also specified some structural constraints, namely that each pattern should start at the beginning of the measure (first character should be a |) or that it should start at the first l before a measure bar. That is, patterns should have one of the following two structures:

- $| < pattern >$ or
- $l * | < pattern >$, where $*$ is a sequence, possibly empty, of only minuses $(-)$.

Table 3 summarizes the results the obtained for SimilaritySegmenter.

Table 3. Occurrences of patterns in data set (SimilaritySegmenter)

Pattern	Player1	Player2	Pattern	Player1	Player2
P1	*yes*	*yes*	P12	no	no
P2	no	no	P13	no	no
P3	no	no	P14	*yes*	no
P4	no	no	P15	no	*yes*
P5	no	no	P16	no	no
P6	no	no	P17	no	no
P7	no	no	P18	no	no
P8	no	no	P19	no	no
P9	*yes*	*yes*	P20	no	no
P10	no	*yes*	P21	no	no
P11	no	no	/	/	/

Due to the possibility of specifying structural constraints, the results were much easier to evaluate (as compared to FlExPat). One surprising result was that the algorithm was able to find rather long patterns. In song Insensatez performed by Player2, for instance, the algorithm found three occurrences of the following pattern (136 characters long, i.e., 8 measures):

```
|A---P---B-p-+---|A---P---B-p-+---
|A---P---B-p-+---|A---P---B-p-+---
|A---P---B-p-+---|A---P---B-p-+---
|A---P---B-p-+---|A---P---B-p-+---
```

Although it is formed by the same cell (`|A---P---B-p-+---`), it is very significant that the algorithm could find such a long pattern. It may have structural implications: Player2 may have used this sequence during some specific part of the song. Looking closer at the data, we notice that, in this specific case, Player2 used this pattern as the accompaniment of most part of the theme.

The algorithm also found some garbage, such as `|A---+-` or even `|B-`, but the main question that arose during the experiment was due to the algorithm's inherent non-determinism[12]: if a pattern in the catalogue wasn't found is it because it was really not played or is it because the algorithm didn't run long enough? This seems to be case of P10 for Player1, which was found several times by FlExPat.

4 CULTURAL ISSUES

At this point, the reader may be questioning: "So what? What does it all has to do with music and culture?". The answer is pretty simple: "A lot!". Behind all this pattern discovery discussion there are several questions related to music, music making, and musical heritage and culture.

As we stated in the Introduction (Section 1), the guitar accompaniment results from the concatenation of patterns done by the performer himself. Due to its nature, part of the knowledge about the accompaniment construction is subjective, which leads to several open questions, such as:

- Are there rhythmic patterns that are preferred by a certain performer or required for a certain musical style? In which situations and in which frequency do they show up?
- Are there variations of these patterns? Is it possible to group these variations in meaningful way? Which variations (timing, dynamics, etc.) are acceptable within a pattern?
- Is it really the case that everything is a pattern, i.e., are there parts that are not recurrent?
- Is it possible to justify the choice of a pattern in terms of other musical features (melody, harmony, tempo, musical structure, style, etc.)?
- Is it possible to build a dictionary of patterns for a given player? Do different players have different dictionaries?
- Is it possible to build a grammar or a set of rules that is able to describe formally how the patterns are chained and/or the rhythmical transformations done by a given performer? If so, what are the relations between grammars from performers p_1 and p_2?

More general questions can also be posed: To which extent patterns used in Bossa Nova music are different from patterns used in Samba? How different is the Samba today (in terms of patterns and pattern usage) from the Samba in the 1920's or 1930's? The answers to these questions will certainly bring much more understanding of how the musicians bring these genres to life, as well as register how they evolved through the years.

Unfortunately, these questions are very hard to answer. They require a huge analysis effort: Large corpuses of data should be analysed in order that conclusions can be drawn. Note that it is indeed possible to analyse some key pieces and have some hints of how these things are interrelated. But, due to the lack of notation, each performance is unique, and not just because of the deviations a performer does, but also most likely unique in the sequence of patterns he plays. So, choosing which performance should be used is arguably an easy and conclusive task. Thus, the use of computational techniques, specially AI techniques, is a key factor in such an effort, since they allow large amounts of data to be analysed.

To close this section, may we come back to one question raised by the experiment, namely the structural malformation of several patterns found by FlExPat. Instead of a problem, it can mean another thing. Many patterns of the catalogue have a common substructure (i.e., subparts of these patterns are equal). It may be the case that the second measure of patterns whose second measure are equal, are frequently concatenated with patterns whose first measure are similar. So, it may be the case that these concatenations are so typical that they are "promoted" to patterns by the algorithm. This completely changes the analysis focus: Instead of analysing the patterns themselves, it may be much more interesting to analyse how these patterns are interconnected, which, in fact, is the very task of building an accompaniment!

[12] SimilaritySegmenter is an evolutionary algorithm, and it does not perform an exhaustive pattern induction [6].

5 FINAL REMARKS AND FUTURE WORK

This paper described an experiment that dealt with extraction of rhythmic patterns from Bossa Nova songs. Sixteen beat tracked MIDI files, representing the recording of several songs by two different players, were represented as a string and thereafter processed by three different string matching algorithms. The objective was to identify in the data set patterns from a previously acquired catalogue.

First results showed that an exact matcher was not able to cope with the variations of the patterns done by the players. With the use of inexact matchers the results were much more interesting: several patterns from the catalogue could be found in the data set. We also pointed out how experiments of this kind can contribute to issues related to musical heritage and culture.

Of course, there are several points for improvement. One that is particularly important is the representation. We used a very simple representation of the events. Attributes like tempo, structural or harmonic information are not represented. The more attentive reader may have even noticed that the actual duration of the events is not represented at all. We just used the onset information, which turned out to work for this experiment. To further investigate the particularities of the patterns, however, we surely need to represent appropriately the duration of each event.

The data acquisition is an important an non-trivial problem. It is important because if we want to draw relevant and significant conclusions about Bossa Nova and/or Samba, we must have a much more representative data set. And it is non-trivial because there are many factors involved varying from the availability and willingness of certain performer to record for us, to copyright issues of the collected material.

We also have collected some other songs, namely some Sambas recorded by three different players. We plan, as soon as we have prepared the data, to perform a similar experiment on this data set.

Finally, cultural heritage exploitation using computers may be frequently associated with data preservation or data storage and the technological issues behind them, such as indexing for fast retrieval, for instance. Of course, that is very important and indeed necessary, but there are other aspects of cultural heritage that are also relevant. These are much more related to the analysis of the musical data, and with the sort of questions we raised before. That is, the computer should work like an X-ray machine and its result, the "*musical X-ray*", may, through careful observation, reveal the music's innards. So, we hope!

ACKNOWLEDGEMENTS

The first author is supported by a research grant from CNPq, Brazil (Conselho Nacional de Desenvolvimento Científico e Tecnológico). The research by the Austrian Research Institute for Artificial Intelligence was supported by the Austrian *Fonds zur Förderung der Wissenschaftlichen Forschung (FWF)* under grant no. Y99-START.

The Austrian Research Institute for Artificial Intelligence acknowledges financial support by the Austrian Federal Ministries for Education, Science and Culture (BMBWK) and for Transport, Innovation and Technology (BMVIT).

REFERENCES

[1] Robert S. Boyer and J. Strother Moore, 'A fast string searching algorithm', *Communications of the ACM*, **20**(10), 762–772, (1977).

[2] *Songbook: Bossa Nova*, ed., Almir Chediak, volume 1–5, Lumiar Editora, Rio de Janeiro, 1990.

[3] Márcio Dahia, Hugo Santana, Ernesto Trajano, Carlos Sandroni, Geber L. Ramalho, and Giordano Cabral, 'Using patterns to generate rhythmic accompaniment for guitar', in *Proc. of Sound and Music Computing (SMC'04)*, pp. 111–115, (2004).

[4] Simon Dixon, 'Automatic extraction of tempo and beat from expressive performances', *Journal of New Music Research*, **30**(1), 39–58, (2001).

[5] Walter Garcia, *Bim Bom: A Contradição sem Conflitos de João Gilberto*, Editora Guerra e Paz, 1999.

[6] Søren Tjagvad Madsen and Gerhard Widmer, 'Evolutionary search for musical parallelism', in *Proc. of the EvoWorkshops2005*, Lecture Notes in Computer Science 3449, pp. 488–497. Springer Verlag, (30 March–1 April 2005).

[7] T. Owens, *Charlie Parker: Techniques of Improvisation, 2 vols.*, Ph.D. dissertation, Univeristy of Los Angeles California, 1974.

[8] Pierre-Yves Rolland, 'FlExPat: Flexible extraction of sequential patterns', in *Proceedings of the 2001 IEEE International Conference on Data Mining (ICDM'01)*, eds., Nick Cercone, Tsau Young Lin, and Xindong Wu, pp. 481–488. IEEE Computer Society, (2001).

[9] Carlos Sandroni, *Feitiço Decente: Transformações do Samba no Rio de Janeiro (1917–1933)*, Jorge Zahar Editor, Rio de Janeiro, 2001.

[10] Ernesto Trajano, Márcio Dahia, Hugo Santana, and Geber L. Ramalho, 'Automatic discovery of right hand fingering in guitar accompaniment', in *ICMC*, pp. 722–725, (2004).

[11] Gerhard Widmer, 'Applications of machine learning to music research: Empirical investigations into the phenomenon of musical expression', in *Machine Learning, Data Mining and Knowledge Discovery: Methods and Applications*, eds., R. S. Michalski, I. Bratko, and M. Kubat, Wiley & Sons, Chichester (UK), (1998).

[12] Gerhard Widmer, 'Machine discoveries: A few simple, robust local expression principles', *Journal of New Music Research*, **31**(1), 27–50, (2002).