MUME 2016 - The Fourth International Workshop on Musical Metacreation, ISBN #978-0-86491-397-5

Rhythm Spaces

Daniel Gómez-Marín, Sergi Jordà, Perfecto Herrera
Music Technology Group - Universitat Pompeu Fabra
Barcelona, Spain
{daniel.gomez,sergi.jorda,perfecto.herrera}@upf.edu

Abstract

We present and discuss our view on rhythm spaces, as metaphors to visualize and interact with rhythms. We take advantage on existing research on alternative low-dimensional music spaces, music cognition, music interaction and intelligent music agents, to justify our position. We then discuss the specific steps and decisions necessary for building rhythm spaces and we advance some of their potential uses as rhythm pattern allocation, retrieval and generation.

1 Introduction

The evolution of electronic dance music (EDM) is necessarily linked to the advancement in techniques for three day-to-day endeavors of an EDM music producer: generating, controlling and processing drum rhythms. Drum arrangements, as the backbone of EDM, are subject to detailed scrutiny and constant transformation in a search for creative sound layers that can catch the attention of an audience. This inventive production scenario is a fertile ground for research and innovation in sound and music computing and interaction.

Current musical systems allow producers to create drum arrangements in diverse forms, using audio loops, sequencing audio onsets on a timeline or creating symbolic drum sequences for samplers or drum machines. In this paper we will focus on symbolic sequences and in ways to organize them in low-dimensional spaces that can be grasped by the users and that can act as organizational and creative devices. These tools could provide significant advantages over other unidimensional orderings (e.g., drop-down menus) that are mostly used in current tools for drum-loop creation.

Several drumming systems involving generative capabilities have been studied and developed from different perspectives. Systems based on the transformation of the onsets in audio loops[1] and symbolic sequences (Burton 1998) (Kaliakatsos-Papakostas et al. 2012) have been created. A wide range of computation approaches have also been used as genetic algorithms, neural networks, stochastic processes, agents (Pachet 2000) and hybrids of these (Bernardes, Guedes, and Pennycook 2010). In general, these systems exploit computer processing for the generation of new rhythms based in some sort of preconceived knowledge. Regarding their context, some tools are aimed to be general approaches to rhythm generation, agnostic to any drumming tradition and focusing on geometrical reflections on pattern repetition (Milne and Dean 2016); others are reactive to any MIDI input (Aucouturier and Pachet 2005); others are rooted in cultural drumming traditions (Sampaio, Ramalho, and Tedesco 2008) while some are loaded with special knowledge of EDM (Eigenfeldt and Pasquier 2013) (Jordà et al. 2016).

Currently, despite the methodology or production context, there is a lack of music creation tools for compiling, meaningfully organizing and visualizing drum patterns (loops). This challenge of organizing elements of one domain and depicting their relations is recurrent in different scientific and engineering scenarios. Given this, well-defined mathematical techniques have been proposed which convert complex collections of elements in a common domain into an easy to understand geometrical construction. Our domain of interest, which is drum patterns, can take advantage of these techniques.

We envision a tool for music production, as a rhythm space where a producer can make sense of her collection of patterns in a significant way, browse through the space to select patterns for a drum machine and generate new patterns that are somehow a combination of those in the space. The features of this interactive rhythm space would be:

- Low dimensionality for ease of navigation.
- Closeness of similar patterns and separation from different patterns.
- Spatial continuity.
- Navigability and real time retrieval of patterns.
- Generativity (in the sense of producing new patterns, different from the ones in the collection, when navigating through an empty region)
- Open architecture, so new patterns could be added or removed form the space.

The rest of this paper first presents research previously done on rhythm spaces, it then discusses some of the open issues yet to be clarified in order to make the envisioned

[1] Spectrasonics' Stylus Rmx software

space feasible, and concludes presenting some current advances towards making them a reality.

2 Research on Rhythm Spaces

2.1 Studies on Monophonic Rhythms

Desain and Honing have an extensive body of work behind modeling human perception of rhythm from a cognitive perspective. In several papers they use a three dimensional space for visualizing rhythms. Each axis of the space represents one of the three inter-onset intervals (IOI) which exist between the four notes of their rhythm. In this informative space, a rhythmic structure is recognized by its position (Desain and Honing 2003) (Honing and Others 2002). Other studies have explained how repetitive acoustic phenomena create a sensation of pulse, eliciting temporal hierarchies and expectancies in subjects, which are confirmed or challenged again by incoming acoustic phenomena (London 2012). Studies in this area propose measures for predicting the similarity sensation between two monophonic patterns (Cao, Lotstein, and Johnson-Laird 2014) (Johnson-Laird 1991) (Gómez-Marín, Jordà, and Herrera 2015b).

Distance metrics can also be computed for rhythms without looking for any perceptual validity, just by taking advantage of certain mathematic properties and relations. Examples of such distances are the edit or the swap distances (Toussaint 2004). Despite some relations found experimentally between the edit distance and the perception of similarity (Post and Toussaint 2011), these types of metrics make no acknowledgment of any of the essential elements of human cognition of rhythms, such as the acquisition of meter, with the consequent omission for fundamental cognitive concepts as reinforcement or challenge of an induced pulse (Cao, Lotstein, and Johnson-Laird 2014). These type of monophonic distances are useful because of their simple computation and common use in computer music.

Another way to study monophonic rhythms is by extracting absolute values from their structure, and not by comparison with others. Descriptors as syncopation (Song, Pearce, and Harte 2015), density or evenness (Milne and Dean 2016) are simple to understand and easy to compute from symbolic patterns. Syncopation, as a concept, is fundamental for understanding rhythm. However, syncopation, density or evenness as a descriptors, they all have a very low granularity[2]. That is, the same value of syncopation is shared by a large number of monophonic patterns. This fact has a blurring effect as thousands of, for example, 16-step rhythms, generating quite different rhythmic sensations, can be characterized with the same descriptor values. Because of that they are not good candidates for spanning low-dimensional spaces.

From a generative perspective, Forth et al. (Forth, McLean, and Wiggins 2008) (Forth, Wiggins, and McLean 2010) (Forth 2012) develop the conceptual construct of a space as a powerful interaction metaphor which they argue can be suited for musical purposes. Two musical applications are derived from their study: first, a monophonic rhythm space that deals with 16-step rhythmic possibilities, and second, the more common and explored notion of timbre space. Their theoretical depiction of a conceptual spaces is thoroughly developed and their rhythm space is modeled following London's approach to rhythm (London 2012).Their tools are not validated as generation systems but by successfully classifying symbolic musical pieces in different ballroom dance music genres.

2.2 Studies on Polyphonic Rhythms

Monophonic and polyphonic rhythm spaces for small sets of rhythms have been created by Gabrielson as cited in (Dowling and Tighe 2014). His polyphonic rhythm spaces organize prototypical rhythms of different musical styles in three dimensions (Gabrielsson 1973). The construction of these spaces is based on similarity ratings given by listeners that compared pairs of patterns with the subsequent application of multi-dimensional scaling (MDS) to obtain the space. MDS is a dimensional reduction technique, commonly used as a tool to visualize multi-dimensional arrangements of data in few dimensions while preserving, as much as possible, the local similarities derived from the raw ratings (Shoben and Ross 1987). Meanings for the space axes are proposed by analyzing the three dimensional arrangement and trying to explain the position of the rhythms in each axis. According to Gabrielson, the main three features for discriminating rhythms found in his polyphonic spaces are (1) meter, (2) differences in basic pattern and (3) uniformity versus variation. Other studies explore the factors which influence polyphonic similarity (Witek et al. 2014). They argue that the low-pitch sounds in a polyphonic rhythm have the highest effects in human perception of syncopation. Other researches try to explain polyphonic similarity as the sum of monophonic similarity (Gómez-Marín, Jordà, and Herrera 2015a). At this moment, we are developing cognitive-inspired metrics to measure the similarity between polyphonic patterns with promising results. Some advances, without any aim to capture perceptual validity, have also been reported (Sampaio, Ramalho, and Tedesco 2008).

There are some other authors that have dealt with rhythm spaces in a polyphonic music audio (retrieval) context. Rhythm spaces are implicit in many MIR studies involving rhythm descriptors (Ellis and Arroyo 2004) (Rocamora, Jure, and Biscainho 2014). Here, spaces are rarely explicitly depicted or used as such, probably because the multi dimensionality of the constructions and because the aims lean more towards classification (Chen and Chen 1998).

2.3 Dimensional Reduction

The MDS technique mentioned in the previous section is a subset of existing dimensional reduction techniques (DRT) such as principal components analysis (PCA). These methodologies are a means, in a variety of contexts, to organize and visualize data. Commonly, visualization of data is based on low dimensional structures (2D or 3D) while the information displayed can convey more dimensions by resourceful use of color, form and symbols (Keim 2002). In cognitive science, as briefly described by Gabrielson's research (Gabrielsson 1973), a low dimensional cognitive

[2] The different syncopation values that can be measured from all possible 16 step patterns are very few compared to the amount of patterns that can be created.

space obtained by DRT can reveal latent relevant characteristics of how a domain is understood by humans, which is not evident in their similarity judgments and perhaps, unknown by them. The value of MDS and any other DRT is preserving the high dimensional relations between the studied stimuli in the resulting low dimensional space. The methodology for creating a space using MDS in a given domain is clear. The starting point is to define a set of stimuli from the domain. Then, subjects evaluate the stimuli pairwise, their results are unified and assigned to each pair in the set, thus obtaining a similarity matrix. Then, MDS is used on the matrix, specifying the desired dimensionality of the expected resulting space. Finally, the result is a set of coordinates for each element on the stimulus set. This methodology is widespread in cognitive sciences and is the foundation of contemporary understanding of many domains such as color (Shepard 1962), timbre (Grey 1977), pitch (Krumhansl 1979) or tactile textures (Hollins et al. 2000).

3 Constructing Rhythm Spaces

3.1 A Meaningful Rhythm Space

Different DRT can be used to simplify relations between elements in a collection so they can be represented in a low dimensional space. However, when the relations are based on human similarity judgments, the resulting structure might go beyond geometry and model a mental representation. While these spaces might not completely account for a theoretical model, they are a reasonable arrangement of how a cognitive model might be (Shoben and Ross 1987).

The importance of these spaces, besides being helpful to gain insight of mental processes, is that they can be considered stable and as such be useful to make predictions on human behavior (Shepard 1987), for example to predict how dissimilar a pair or elements in the space might be perceived. As the spaces are specific geometrical constructions, mathematics apply throughout them, and relations among the elements are based on structured principles such as location, distance, magnitude, direction and area, to name some. Predictions in these cognitive spaces, and relations among their elements in general, can be quantified and systematized, opening a door to systems which make use of these spaces in a human meaningful way. A simple but illustrative example of the power of cognitive spaces is the possibility to obtain a structure of rhythms organized by zones or regions. These can be explored locally to retrieve specific types of patterns or can be traveled establishing rhythmic progressions in different trajectories. What was an amorphous collection of elements is then converted to an organized structure analogous to a mental representation of it.

3.2 Two Approaches for Constructing Spaces

The problem of creating a rhythm space could be faced from a straight cognitive perspective. This would imply the application of the methodology exposed above: selecting prototypical EDM patterns, getting similarity ratings, applying MDS and (optionally) determining some acoustic-musical features that correlate with the axes. Given that the data can be represented in an euclidean space, the axis should reveal

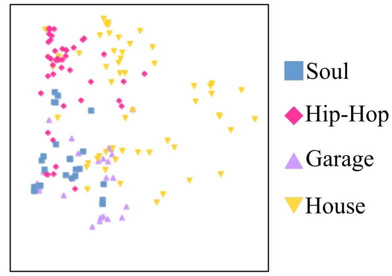

Figure 1: A rhythm space obtained with one of our similarity metrics and a collection of patterns in different styles.

the internal features that influence our perception of similarity. Ideally, by following this procedure, the properties of polyphonic rhythms that guide our similarity judgments could be determined. Provided this information could be deduced, these properties could then be modeled and extracted from symbolic rhythms. However, the domain of all polyphonic EDM rhythms is very wide and the number of pairs to be evaluated is larger. Given that evaluation is a time consuming activity which should be completed by different subjects, the amount of time to be invested in the creation of this space perhaps limits its feasibility.

Another approach to the problem could be based in the research being done in polyphonic similarity of drum patterns. This would require to implement the findings on polyphonic perception (Cao, Lotstein, and Johnson-Laird 2014) and contrasting them with the polyphonic distances reported elsewhere (Gómez-Marín, Jordà, and Herrera 2015a) and with others as proposed in CinBalada (Sampaio, Ramalho, and Tedesco 2008). We expect meaningful similarity predictions could be computed for a set of polyphonic patterns and that some spaces could be built using this information.

We therefore propose a methodology for constructing rhythm spaces along this line. Our approach consists of replacing subject-based similarity ratings by automated computation of similarity distances and checking the perceptual appropriateness as ”a posteriori” operation. In our methodology, a set of polyphonic patterns is input and polyphonic distance metrics are automatically computed. The similarity matrix obtained is then be processed by a DRT, and set to output a low dimensional space. The result is a space where all patterns are located preserving the relations computed by the similarity metrics.

We are currently evaluating methods to predict similarity judgments. With the partial results we have gathered, we have created different algorithms to compute polyphonic similarity between drum patterns. These in turn have been used with different collections of drum patterns in different styles. This exercise has helped us exploring the effect of our metrics in the creation of polyphonic rhythm spaces. Despite the metrics being under development, good spaces have been created. It can be observed in Figure 1 how musical relations between musical styles come through when creating spaces based on drum patterns of different styles.

3.3 Properties of a Rhythm Space

Following our methodology presented on the previous section, the desired features we discussed in the introductory section could be met. The resulting spaces would be low dimensional. As mentioned, the resulting number of dimensions can be predefined for any of the DRT. New spaces can be computed automatically so adding a new pattern would be simple. Having the exact coordinates of each element, an interactive system for navigating could be easily implemented. This structure can be searched using coordinates and the patterns can then be retrieved by proximity to a pointer: when the pointer is over a pattern in the space, the pattern would be retrieved.

3.4 A Generative Rhythm Space

A discrete low dimensional space composed of known patterns can be explored continuously, and new patterns can be retrieved whenever exploring an empty region. Converting a discontinuous collection of samples into a diaphanous and smooth space is indeed a generative action, which implies the prediction of the output in positions that have not been sampled. Currently we have developed some algorithms, to achieve continuous pattern generation.

Our algorithms expand the notion of two-track blending, as typicaly performed by DJs. A smooth one-stage transition between two different rhythms A and B suggests that a third pattern C is created. C must have features that resemble A but also B. Figure 2 presents a progressive interpolation between two patterns based on one of our algorithms. In this example patterns A and B have three onsets in common: a kick on step 1 a snare on step 5 and 13. At three equidistant values of the interpolation between A and B snapshots of the resulting interpolated rhythm are presented. Our algorithm takes care of introducing onsets at each step based on the interpolation value. Figure 2 shows how pattern B onsets are progressively introduced as the interpolation value goes from A to B.

Expanding this notion, we have created algorithms where three patterns are interpolated and a fourth resulting pattern extracted. At every stage of the navigation through the 2D rhythm space, the position of a pointer, controlled by a user, is always inscribed inside the area of a triangle. Based on the drum patterns located on each vertex of the triangle and the distance from the pointer to each pattern, our algorithms generate a new pattern as an output.

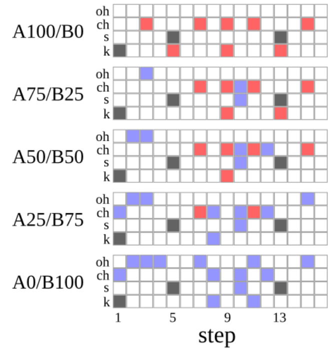

Figure 2: Interpolation between pattern A (top) and pattern B (Bottom). Onset colors are red for Pattern A, blue for Pattern B and gray common onsets. Instruments are kick (k), snare (s), closed hi-hat(ch) and open hi-hat (oh)

4 Conclusions

Automatic creation of drum spaces is possible from a mathematical perspective. Given this, there could be two different resulting scenarios by taking the approach described. On the best case scenario, if the metrics developed are really good predictors of polyphonic similarity, the spaces will inherit the essential relations computed by the distances. This would imply a coherent ordering of the elements in the space. These spaces would also reveal orderings of the elements along the axes, as independent drivers of similarity perception and zones of similar patterns stylistic relations. In the case where metrics for polyphonic similarity are not completely precise in their predictions, the resulting spaces would still be useful for navigation. Users could take advantage of the distribution of patterns in the low-dimensional space as a means for exploring a collection, as some regions might reveal themselves as useful and valid. Both scenarios propose a fluid and novel means of rhythmic interaction. The generative capabilities of the space would account as a powerful tool for music producers, expanding current systems and also providing a means for musical creativity.

A crucial factor that could definitely contribute to the impact of this tool is the visual display of the information, as graphical elements can affect positively the comprehension and use of the tool. Arriving at a final version of this generative space requires human evaluations in different aspects. Experiments must be carried out for determining the factors that influence rhythmic similarity perception. Resulting spaces must studied from a user experience framework. Diverse approximations to rhythm interpolation must also be studied to understand their impact in music production scenarios.

5 Acknowledgments

This research has been partially supported by the EU funded GiantSteps project[3] (FP7-ICT-2013-10 Grant agreement nr 610591).

References

Aucouturier, J.-J., and Pachet, F. 2005. Ringomatic: A real-time interactive drummer using constraint-satisfaction and drum sound descriptors. In *Proc. of the International Conf. on Music Information Retrieval*, 412–419.

Bernardes, G.; Guedes, C.; and Pennycook, B. 2010. Style Emulation of Drum Patterns By Means of Evolutionary Methods and Sta- Tistical Analysis. In *Proc. of the Sound and Music Conf.*, 1–4.

Burton, A. R. 1998. *A Hybrid neuro-Genetic pattern evolu-*

[3]http://www.giantsteps-project.eu

tion system applied to musical composition. Ph.D. Dissertation, University of Surrey.

Cao, E.; Lotstein, M.; and Johnson-Laird, P. N. 2014. similarity and families of musical rhythms. *Music Perception: An Interdisciplinary Journal* 31(5):444–469.

Chen, J. C. C., and Chen, A. L. P. 1998. Query by rhythm: An approach for song retrieval in music databases. In *Research Issues In Data Engineering, 1998.'Continuous-Media Databases and Applications'. Proc., Eighth International Workshop on*, 139–146. IEEE.

Desain, P., and Honing, H. 2003. The formation of rhythmic categories and metric priming. *Perception* 32(3):341–365.

Dowling, W. J., and Tighe, T. J. 2014. *Psychology and music: The understanding of melody and rhythm*. Psychology Press.

Eigenfeldt, A., and Pasquier, P. 2013. Evolving structures for Electronic Dance Music. In *Gecco'13: Proc. of the 2013 Genetic and Evolutionary Computation Conf.*, 319–326.

Ellis, D. P. W., and Arroyo, J. 2004. Eigenrhythms: Drum pattern basis sets for classification and generation. In *ISMIR 2004: 5th International Conf. on Music Information Retrieval*, 101–106.

Forth, J.; McLean, A.; and Wiggins, G. 2008. Musical creativity on the conceptual level. In *Fifth International Joint Workshop on Computational Creativity, Madrid*, 17–19.

Forth, J.; Wiggins, G. A.; and McLean, A. 2010. Unifying conceptual spaces: Concept formation in musical creative systems. *Minds and Machines* 20(4):503–532.

Forth, J. 2012. *Cognitively-motivated geometric methods of pattern discovery and models of similarity in music*. Ph.D. Dissertation, Goldsmiths, University of London.

Gabrielsson, A. 1973. Similarity ratings and dimension analyses of auditory rhythm patterns: II. *Scandinavian Journal of Psychology* 14(3):161–176.

Gómez-Marín, D.; Jordà, S.; and Herrera, P. 2015a. Pad and Sad: Two awareness-Weighted rhythmic similarity distances. In *16th International Society for Music Information Retrieval Conf. ISMIR.*

Gómez-Marín, D.; Jordà, S.; and Herrera, P. 2015b. Strictly Rhythm: Exploring the effects of identical regions and meter induction in rhythmic similarity perception. In *11th International Symposium on Computer Music Multidisciplinary Research (CMMR), Plymouth*.

Grey, J. M. 1977. Multidimensional perceptual scaling of musical timbres. *The Journal of the Acoustical Society of America* 61(5):1270–1277.

Hollins, M.; Bensmaïa, S.; Karlof, K.; and Young, F. 2000. Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling. *Perception & Psychophysics* 62(8):1534–1544.

Honing, H., and Others. 2002. Structure and interpretation of rhythm and timing. *Tijdschrift voor Muziektheorie* 7(3):227–232.

Johnson-Laird, P. N. 1991. Rhythm and meter: A theory at the computational level. *Psychomusicology: A Journal of Research in Music Cognition* 10(2):88–106.

Jordà, S.; Gómez-Marín, D.; Faraldo, Á.; and Herrera, P. 2016. Drumming with style: From user needs to a working prototype. In *Proc. of the New Interfaces For Musical Expression conf.*, in print.

Kaliakatsos-Papakostas, M. A.; Floros, A.; Kanellopoulos, N.; and Vrahatis, M. N. 2012. Genetic evolution of L and FL-systems for the production of rhythmic sequences. In *Proc. of the 14th annual conf. companion on Genetic and evolutionary computation*, 461–468. ACM.

Keim, D. A. 2002. Information visualization and visual data mining. *Visualization and Computer Graphics, IEEE Transactions on* 8(1):1–8.

Krumhansl, C. L. 1979. The psychological representation of musical pitch in a tonal context. *Cognitive Psychology* 11(3):346–374.

London, J. 2012. *Hearing in time: Psychological aspects of musical meter*. Oxford University Press.

Milne, A. J., and Dean, R. T. 2016. Computational creation and morphing of multilevel rhythms by control of evenness. *Computer Music Journal* 40(1):35–53.

Pachet, F. 2000. Rhythms as emerging structures. In *Proc. of 2000 International Computer Music Conf., Berlin, ICMA*.

Post, O., and Toussaint, G. 2011. The edit distance as a measure of perceived rhythmic similarity. *Empirical Musicology Review* 6(3):164–179.

Rocamora, M.; Jure, L.; and Biscainho, L. W. P. 2014. Tools for detection and classification of piano drum patterns from candombe recordings. In *Proc. of the 9th Conf. on Interdisciplinary Musicology (CIM14), Berlin, Germany*, 382–387.

Sampaio, P. A.; Ramalho, G.; and Tedesco, P. 2008. CinBalada: a multiagent rhythm factory. *Journal of the Brazilian Computer Society* 14(3):31–49.

Shepard, R. N. 1962. The analysis of proximities: Multidimensional scaling with an unknown distance function. I. *Psychometrika* 27(2):125–140.

Shepard, R. N. 1987. Toward a universal law of generalization for psychological science. *Science* 237(4820):1317–1323.

Shoben, E. J., and Ross, B. H. 1987. Structure and process in cognitive psychology using multidimensional scaling and related techniques. *The Influence of Cognitive Psychology on Testing*.

Song, C.; Pearce, M.; and Harte, C. 2015. SYNPY: a python toolkit for syncopation modelling. *Proc. of the 12th Int. Conf. on Sound and Music Computing (SMC-15)* 295–300.

Toussaint, G. T. 2004. A Comparison of Rhythmic Similarity Measures. In *Proc. of the 7th International Society for Music Information Retrieval Conf. (ISMIR 2004)ISMIR*, 242–245.

Witek, M. A. G.; Clarke, E. F.; Kringelbach, M. L.; and Vuust, P. 2014. Effects of polyphonic context, instrumentation, and metrical location on syncopation in music. *Music Perception: An Interdisciplinary Journal* 32(2):201 – 217.