MAPPING TIMING STRATEGIES IN DRUM PERFORMANCE

George Sioros **Guilherme Schmidt Câmara** **Anne Danielsen**
RITMO Centre for Interdisciplinary Studies in Rhythm, Time and Motion
Department of Musicology, University of Oslo
{georgios.sioros, g.s.camara, anne.danielsen}@imv.uio.no

ABSTRACT

How do drummers express different timing styles? We conducted an experiment in which we asked twenty-two professional drummers to perform a simple rhythmic pattern while listening to a metronome. Here, we investigate the strategies they employed to express three different instructed timing profiles for the same pattern: "On", "Pushed" and Laid-back. Our analysis of the recordings follows three stages. First, we compute sixteen boolean features that capture the microtiming relations of the kick, snare and hi-hat drum onsets, between each other and with regards to the metrical grid. Second, we construct a microtiming profile (mtP) for every performance by averaging the boolean features across the recording. An mtP codifies the frequency with which the various features were found in a performance. Third, through a "similarity profiles" hierarchical clustering analysis, we identify groups of recordings with significant similarities in their mtPs. We found distinct strategies to express each intended timing profile that employ specific combinations of relations between the instruments and with regards to the meter. Finally, we created a map that summarizes the main characteristics of the strategies and their relations using a phylogenetic tree visualization.

1. INTRODUCTION

In groove performance, it has been assumed that musicians can apply different timing 'feels' to a given pattern by, amongst other things, subtly altering the temporal location of events at the 'micro-rhythmic' level by playing either slightly early ('pushed') or late ('laid-back') in relation to other players' rhythm, a metronomic beat reference or simply their own internal pulse [1, 3, 6, 7, 9, 10, 15, 19]. Typical reported values of microtiming deviations in performance range from 0 ms (no displacement) to 50 ms or more, depending on instrument, tempo and genre [2, 11, 13, 22] An instructed timing experiment by Danielsen et al. [9] showed that drummers were able to consistently play a snare-drum pattern with laid-back and pushed feel significantly behind- and ahead-of an instructed on-beat performance, respectively, with similar values. In polyphonic drumkit performance, expert drummers are able to control the degree of onset timing asynchrony between the various constituent drum instruments. These inter-instrument onset asynchronies may play a role in the production and perception of groove timing feel, since both magnitude and order of onset asynchrony between near-simultaneous events have been previously shown to affect judgements of timing in perceptual experiments with musical stimuli [12, 14, 25].

In order to explore potential interactions between instructed timing feel and various audio/motion features in drum-kit performance, a series of experiments was conducted by Câmara et al. [4] where participants played a simple 'back-beat' pattern with On-beat, Pushed and Laid-back timing feel along to a metronome. In the present study, we analyze the data from one of these experiments, limiting our focus towards investigating the extent to which professional drummers employed different strategies in order to achieve the instructed timing feel in terms of the magnitude and order of onset asynchrony between the instruments of the drum-kit themselves, as well as in relation to a metrical reference grid. We hypothesize that drummers chose different elements (read: instruments) of the rhythmic pattern to produce in sync, late and early timing performance for the On, Laid-back and Pushed timing condition, respectively. For example, in order to achieve the same timing instruction, one group of participants may have focused on the relation between two drum instruments, where one led and the other followed, while another group instead on the relation between both instruments and the metrical grid, and yet another group may have incorporated a combination of two such approaches. In other words, for the same timing instruction, drummers may produce different combinations of microtiming onset strategies in order to communicate the same intended timing feel. This article focuses on the novel analysis we developed which aims at mapping and identifying these potentially different strategies.

At the core of our analysis lie the microtiming profiles – structures that effectively codify the onset asynchronies of the instruments as the probabilities or frequencies with which they occur in the performances of the participants. A hierarchical classification of the performances based on their microtiming profiles reveals specific timing strategies, which are summarized as microtiming archetypes that capture the main characteristics of the clusters in a symbolic form. Finally, a visualization of the clustering result as a phylogenetic tree enables us to better understand and identify these strategies.

The rest of the article is divided in four sections. In section 2, we describe the experiment. In section 3, we describe the analysis method, the microtiming profiles in 3.1 and their clustering in 3.2. In section 4, we present the results of our analysis. In the final section 5, we discuss methodological issues.

2. EXPERIMENT

22 male drummers, 22-64 years of age [M = 36, SD = 11] participated in the experiment. All of them were active part-time or full-time musicians and had between 4 and 40 years of professional performance experience [M = 16, SD = 11]. All were familiar with at least one groove-based performance tradition, typically either jazz, funk/soul/R&B, hip-hop, rock, or reggae. Two participants' data were excluded from the analysis: one due to technical issues during the recording process, and the other was deemed to not have successfully understood the task based on responses from a follow up interview.

The participants were instructed to play a standard 'back-beat' pattern (see) ubiquitous in groove music and highly familiar to drummers. They performed along to a metronome (woodblock) track at a tempo deemed comfortable in a pilot of the experiment (96 b.p.m) in 3 different timing style conditions:

1. in a *laid-back* manner, i.e. behind-the-beat (condition: Laid-back)
2. in a *pushed* manner, i.e. ahead-of-the-beat (condition: Pushed)
3. in an o*n-the-beat* manner, (condition: On)

At the beginning of the experiment, a practice round was given to in order to allow for participants to accustom themselves to the following timing style conditions ('Laid-back', 'On', and 'Pushed'), which were subsequently randomized. Each timing condition trial lasted for approximately 70 seconds where participants began to play as soon as they had entrained with the timing reference track. This resulted in approximately 200 hi-hat and 50 snare and kick drum strokes captured per trial.

For our drum instrumental setup, we used the following equipment: a Gretsch acoustic metal snare drum (Gretsch Drums, CT), 7 in. deep, 14 in. wide, with a Remo Emperor X drumhead (Remo, CA) with a thin plastic muffle ring; a Gretsch 21-in. bass drum with Remo FA batter drumhead; a Pearl hi-hat stand with 14" Yamaha cymbals.

Pilot tests of the sound recordings revealed that close-microphone techniques with dynamic microphones led to too much leakage between the different drum signals, therefore AKG C411 contact microphones (AKG, Austria) were used instead and placed on the top skins of the kick and snare, and on the top cymbal of the hi-hat.

Figure 1: Standard back beat grove pattern in 4/4 meter. Upper notes in the score denote hi-hat cymbal; the middle notes, the snare drum; the bottom notes, the kick drum

3. ANALYSIS

Since the focus of this investigation is the microtiming relations of the drum instruments, we create microtiming profiles (mtPs) of the performances for all participants and instructed timing condition trials, comprised of set of features that capture those relations. Based on how similar the mtPs are, we group them using a hierarchical clustering algorithm and construct archetypes that summarize the main characteristics of each group. Finally, we map the relations between the groups' recordings using a phylogenetic tree visualization.

Our analysis and all following computations are based on the temporal location of onsets of individual strokes from each instructed timing condition recording that were calculated using an adaptation of an existing onset detection algorithm of the MIRtoolbox [17] which will be detailed in a forthcoming publication of our group.

We describe the mtPs and the measurements we use to obtain them in subsection 3.1, then present the clustering results and their visualization in subsection 3.2.

3.1 Microtiming Profiles

To obtain the mtP of a recording we first extract a set of sixteen boolean features that capture the microtiming relations between the strokes of the snare, kick drum and hi-hat cymbals relative to each other, as well as to the location of the metrical grid. The boolean features are calculated for each 4/4 measure of a timing condition trial while the mtP is calculated as the average of the boolean features across all measures of a trial. The mtPs were inspired by the motion templates designed by Müller and Röder [20] to describe geometric relations of the human body for the purpose of the analysis of body movements.

Kick on beat 1	x2 features per instrument late/early relative to the hi-hat cymbal
Snare on beat 2	
Kick on beat 3	
Snare on beat 4	
Kick + Hi-hat on beat 1	x2 features per beat both instruments late/early relative to the metrical grid
Snare + Hi-hat on beat 2	
Kick + Hi-hat on beat 3	
Snare + Hi-hat on beat 4	

Table 1: Summary of the sixteen boolean features extracted from the recordings for each bar.

Each feature tests whether an instrument is late or early with respect to a certain reference time point. The first eight features use the onsets of the hi-hat strokes as a reference and the other eight features use the metrical grid as a reference. For instance, feature 1 tests whether the kick drum follows the corresponding hi-hat cymbal, while feature 9 tests whether both the kick and hi-hat occur after the position of the respective beat. Table 1 summarizes the sixteen binary features extracted from the recordings of each trial.

The above features depend on "tolerance" thresholds with which two instrument's strokes are considered synchronous with each other (for features 1 to 8) or with which the instruments' strokes are considered to occur late or early relative to the beat of the metrical grid (for features 9 to 16). For instance, when the inter-onset interval (IOI) between a pair of kick and hi-hat strokes is greater than the respective synchronicity threshold, feature 1 (kick later than hi-hat) is *true*. In the opposite case where IOI is sub-threshold, both features 1 and 5 (kick later and earlier than hi-hat) are *false*, since the pair is considered to be synchronous.

Furthermore, to determine the relation between the three drum strokes and the metrical grid with which drummers used as a beat reference, we need first to determine the location of this grid. Although one might intuitively assume that the onset of the sounding metronome would correspond to that location, it has been repeatedly observed that people tend to tap to a steady pulse systematically earlier than the actual pulse—a phenomenon known as negative mean asynchrony (NMA) [8]. As such, it may be assumed that the internal pulse scheme with which drummers operate with, that is, their subjective metrical grid, is slightly anticipated.

Even though percussionists and drummers tend to display lower NMA than other musicians in both in-phase synchronous tapping [21] and drumming [11] experiments, NMAs still tends to vary significantly between individuals [8]. Therefore, it is difficult to assume a single global NMA value for all the drummers. Similarly, the two thresholds values described above cannot easily be set universally. In what follows, we will describe how we obtain individual values for these parameters for each drummer based on the performance of their On timing condition, essentially turning the On recordings into a baseline reference. We will discuss the reasoning behind this choice as well as some of its implications in section 5.

For the synchronicity threshold values used in features 1-8, i.e. the tolerance with which two coinciding drum strokes are considered synchronous or not, we use the variability of the IOI between the hi-hat and coinciding kick or snare strokes on each of the four main quarter-note beats of the 4/4 metre (kick + hi-hat on beats 1 and 3, snare + hi-hat on beats 2 and 4). For each drummer, we first calculate the standard deviation of the inter-instrument IOIs of the corresponding beats in each measure of their respective On condition recording. This yields four separate values, two for each hi-hat + kick, and hi-hat + snare, feature. To be conservative, the synchronicity threshold for a drummer is chosen as the maximum of these four values, then multiplied by 2. Thus, a kick or snare stroke is considered to occur as either asynchronously late or early relative to its coinciding hi-hat stroke when their onsets satisfy the following inequalities:

Late: $Onset(i,j) - HiHat(i,j) > SyncThr(i)$

Early: $Onset(i,j) - HiHat(i,j) < -SyncThr(i)$

$$SyncThr(i) = 2 \times \max_{k=1}^{4}\{STD_{ON}^{i}(IOI\ at\ beat\ k\}$$

where j is the beat number from a recording of drummer i.

The synchronicity threshold used in features 9 – 16 on the other hand, i.e. the tolerance with which a stroke is considered to occur late or early relative to a corresponding subjective metrical beat, is based on the timing variability of the hi-hat strokes in the On condition. The hi-hat is chosen because it can be considered more of a 'time-keeper' instrument than the other drums, thus serving more aptly as a proxy for the beats of the drummers' subjective metrical reference. To account for the commonly observed anticipation of the beats of the metronome (henceforth abbreviated as AoB), we first calculate the mean position of the hi-hat strokes relative to the metronome in the On recordings for each drummer. 18 out of 20 drummers displayed NMA of hi-hat strokes relative to the actual metronome, consequently yielding negative AoB values. Two participants displayed either no NMA or minutely positive mean asynchrony, and in these cases the AoB was set to 0 (no anticipation). Finally, we consider any drum stroke as occurring asynchronously late or early in relation to the beats of the metrical grid according to the following inequalities:

Late: $Onset(i,j) - beat(j) - AoB(i) > 2 \times STD_{ON}^{i}(HiHat)$

Early: $Onset(i,j) - beat(j) - AoB(i) < -2 \times STD_{ON}^{i}(HiHat)$

where j is a drum stroke of drummer i, $beat(j)$ is the corresponding position of the metronome, and $STD_{ON}^{i}(HiHat)$ is the standard deviation of the IOI of the hi-hat from the respective metrical beat positions in the On recording of the same drummer. For one of the features 9-16 to be *true*, the onsets of both strokes—hi-hat and kick or snare—must be either early or late. In all other cases, including when one stroke is early and the other late, the corresponding features would be *false*.

Two main observations must be made about the boolean features. First, all boolean features form mutually exclusive pairs. For instance, feature 9 (strokes on beat 1 are early) and feature 13 (strokes on beat 1 are late) cannot be both *true* for the same bar of a recording. However, they can both be *false*, in which case the combination of the two strokes is considered on the beat. Second, features 1-8 (kick and snare relative to hi-hat) and features 9-16 (onsets relative to metrical grid) are independent. For instance, a kick onset can be early relative to the respective hi-hat onset (feature 5 *true*) while at the same time they are both late relative to the beat position (feature 9 *true*).

The final microtiming profiles (mtPs) of the recordings are calculated by averaging the boolean features across each timing condition recording for all drummers. As the boolean features take either *true* (1) or *false* (0) values, averaging them results in values in the range [0, 1], which represent the frequency with which a feature was

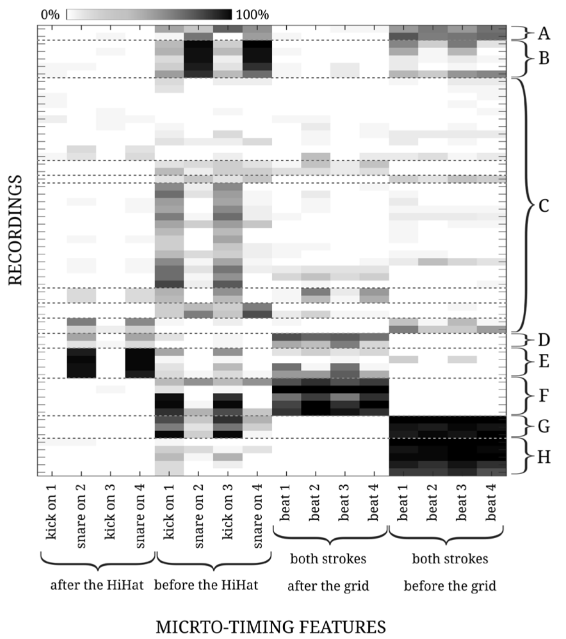

Figure 2: Microtiming profiles (mtPs) for all recordings shown as a greyscale image representing the probability or frequency with which a feature is encountered in a recording. Features are laid along the horizontal axis. On the vertical axis, recordings are sorted and grouped based on the proximity of the SIMPROF clusters (see section 3.2 and Figure 3). Horizontal dashed lines represent the cluster boundaries. The corresponding group mt archetypes are marked on the right with the letters A-H.

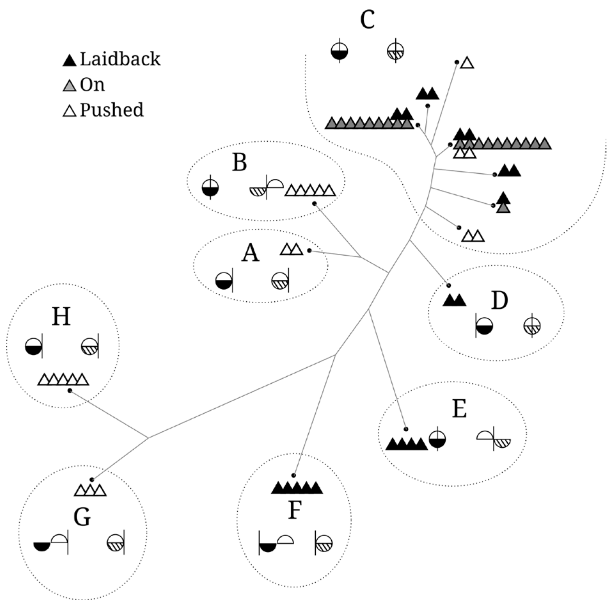

Figure 3: Hierarchical clustering presented as a phylogenetic tree (unrooted, equal daylight visualization). Each triangle corresponds to the microtiming profile of a single recording. Letters A-H are used to label the clusters. Next to each cluster the corresponding mt-Archetype is shown. Label C is assigned to a group of proximal clusters which correspond to the same mt archetype.

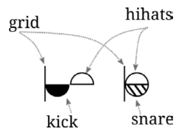

Figure 4: Explanation of an mt-Archetype symbol.

encountered in a recording. The mtPs of all the recordings are visualized in matrix form in Figure 2.

3.2 Hierarchical clustering

We sought to identify the extent to which drummers implemented distinct microtiming strategies for different timing conditions and whether they formed different groups. To this end, we used an agglomerative, hierarchical cluster analysis of the mtPs. A hierarchical clustering was preferred to other clustering methods, like k-means, since it is flexible in that it does not require an a-priori number of clusters to be determined, nor does it impose restrictions on the distribution of the data. Its only requirement is a similarity metric between the data points. We treated the mtPs as arrays of variables, where the similarity between two mtPs is the Euclidian distance between them.

Hierarchical agglomerative algorithms result in dendrograms by successively joining neighboring data points or groups of previously joined points. A linkage criterion determines the distance between groups of points as a function of the pairwise distances of the points themselves. In this study, we used the common Unweighted Group Average linkage (UPGMA) [18, p. 352].

To create clusters of similar data points, one generally "cuts" the dendrogram at different heights. Here, we borrowed methods from the fields of bioinformatics andecology to identify clusters of mtPs. We used the similarity profiles (SIMPROF) method [5]—as implemented in the Fathom Toolbox for Matlab [16]—to test the statistical significance of the branches' internal structure. SIMPROF takes the form of a series of permutation tests. Beginning at the top of a precalculated hierarchy, thesetests stop the ever finer partitioning into subgroups. When a branch in the hierarchy is deemed to have no internal structure and is therefore homogenous, it is no longer subdivided. Thus, a cluster is formed that is comprised of recordings with "exchangeable" features.

For the permutation test, we set the number of iterations to 1000 and the significance level alpha to 0.05—the probability value at which the hypothesis of an internal structure is rejected. We used the Bonferroni correction [18, p. 745] to progressively adjust the probability values for multiple simultaneous tests (see also parameter mc=*true* of the `f_disprof_clust` function of the Fathom toolbox [16]).

The results of the clustering analysis are shown in Figure 2. The dashed horizontal lines cut the mtP matrix into clusters of recordings that show statically significant similarities. In Figure 3, we present the same result as an unrooted phylogenic tree. Recordings of the various Laidback, On and Pushed performances are represented as black, grey or white triangles, respectively, while the distance between the clusters of recordings corresponds to the UPGMA linkage criterion [18, p. 352].

As our aim is to identify distinct timing strategies in the recordings, we summarize the main characteristics of each cluster into microtiming archetypes. An archetype is computed by first averaging the probabilities of the features that belong in beats 1 and 3 (where the kick strokes occur) as well as the ones that belong in beats 2 and 4 (where the snare strokes occur). The two groups of probabilities describe the relation of the kick/hi-hat and snare/hi-hat combinations of strokes between them and in relation to the meter. Those relationships are reduced into archetypes according to whether the average probability of each feature is above or below 50%. An example of such an archetype is shown in Figure 4. In this example, the average probability of both a kick and corresponding hi-hat stroke to be late relative to the beat is above 50%. At the same time, the average probability of a kick stroke being ahead of the corresponding hi-hat stroke is also above 50%. In contrast, the snare strokes have a probability below 50% to either occur ahead or after the corresponding hi-hat stroke. In other words, most of the snare strokes are considered as synchronous with their corresponding hi-hat strokes.

Clusters with common microtiming archetypes that are relatively proximal in the phylogenetic tree are grouped together to create an overall map of strategies. In Figure 2, these groups are labeled with the letters A to H and in Figure 3 the same groups are annotated with their corresponding archetype.

4. RESULTS

The classification of the mtPs shows that the majority of the Laid-back and Pushed performances form separate homogenous clusters. The purely Laid-back clusters are characterized by generally late timing while the Pushed ones by early timing, as expected. However, the analysis reveals that drummers implemented distinct strategies to express a given timing feel by focusing on different rhythmic elements. The microtiming archetypes assigned to the various group clusters highlight those elements.

More specifically, there are 5 purely Laid-back clusters which comprise 15 out of the 20 Laid-back performances. However, 4 of those recordings (split into 2 clusters in group C) are proximal to the two On clusters and are thus subsumed by the very same archetype. The other 11 are split in three clusters each forming a separate group (D, E and F). All three groups are characterized by late strokes. On the one hand, in group E, the snare stroke is late in relation to the hi-hat and in F, both hi-hat and snare are played late relative to the metrical beat. On the other hand, in group F, the combined kick/hi-hat strokes are late relative to the beat, while the kick additionally precedes the hi-hat.

The Pushed performances are mainly found in purely Pushed clusters (18 out of the 20). However, in similar fashion to the Laid-back condition, a small portion (3) are found proximal to the On clusters in group C. 15 of the remaining Pushed performances form 4 distinct groups (A, B, G and H). All of them are characterized by the early timing of the snare strokes. In group B, the snare precedes the hi-hat although the snare/hi-hat combination is considered as synchronous with the beat. Group B exhibits the inverse pattern of the Laid-back group E (snare early in relation to hi-hat). Similarly, the Pushed group G has its counterpart in the Laid-back group F, with both instruments being early in relation to the beat instead of late but the kick still precedes the hi-hat.

In groups A and H the rhythmic pattern appear to be simply shifted early in relation to the beat. However, although the two groups correspond to the same archetype, they are relatively distant on the tree. A closer look at their mtPs in Figure 2 reveals that in group A, the strokes are anticipating the beat significantly less often than in group H. This can also be seen in the proximity of the group A to group C which is dominated by the On recordings. The similar proximity of the Laid-back group D to the On group C reflects the analogous weak late timing features of the mtPs in comparison with the other Laid-back groups (E, F).

The On performances are all found in Group C which is characterized by synchronous on-the-beat stroke onsets. Within the group, the On recordings are split into two clusters. Examining their mtPs, we see their difference stem from the tendency of some musicians to play the kick drum ahead of the hi-hat.

5. DISCUSSION

In this study, we present findings of an experiment in which professional drummers performed the same rhythmic pattern with an On, Laid-back and Pushed timing feel. We found that participants used more than one distinct onset microtiming pattern for each intended timing instruction (see section 4). A more in-depth discussion of the timing strategies and their musicological implications will be undertaken in an upcoming publication. In the present discussion, we will focus on methodological issues concerning, first, the various parameters of the analysis and their implications, and second, the interpretation of the clustering results.

Our analysis begins with the encoding of the onset asynchronies found in the performances of the drummers into sets of boolean features. Microtiming profiles of the performances are calculated then by averaging those features over the respective recordings. The mtPs codify the probability or frequency with which each feature was encountered in a recording. Finally, by clustering the mtPs we discover the recordings with similar features and group them together.

The boolean features, however simple in their definitions, depend on parameters and thresholds which are crucial to the outcome of the analysis and are closely related to the research questions. In this study, we chose to use individual values for each drummer instead of setting global ones across all participants. This decision partly reflects the fact that phenomena such as negative mean asynchrony (NMA) typically varies between individuals. and, at the same time, naturally follows from our research question that seeks to identify individual strategies that musicians employ. The three parameters used for calculating the boolean features reflect mechanisms relevant to the perception and production of the subtle asynchronies we are studying. Therefore, setting individual values in our analysis corresponds to adopting the 'point of view' of each separate musician independently, whereas global values might instead correspond better to the perception of a 'typical' listener.

The way in which individual values are assigned to each parameter can significantly impact results depending on how the research question is formalized. In our current approach, we chose to derive the individual parameters for each musician based on their respective On performances, essentially rendering the On condition into a baseline from which the other timing conditions were compared against. For instance, whether a stroke is considered as 'on-the-beat' or not depends on whether it was performed late or early relative to the average hi-hat stroke in the On condition. In this case, the research question could perhaps be interpreted instead as "how do musicians differentiate their Laid-back or Pushed from an On timing feel".

Consequently, one might assume that the mtPs of the On recordings contain no meaningful information since, after all, they cannot be different from themselves! Nevertheless, the On performances should not be excluded from the analysis: firstly, their mtPs can still exhibit significant enough differences between the performances to classify them separately (see group C), though those differences can solely be obtained from the onset relations between the instruments themselves, and not with respect to the metrical grid. Secondly, the proximity of the Laid-back and Pushed clusters to the two On clusters in the phylogenetic visualization of Figure 3 is informative insomuch as it is telling of the strong tendency for drummers to differentiate these asynchronous timing feels from the On timing feel.

The further hierarchical clustering of the derived mtPs proved an effective means of identifying several key timing strategies implemented by the drummers. The method groups and sorts the recordings according to their similarity, revealing their relations without the need for a-priori hypotheses about the existence of specific strategies. Although in principle it is possible to analyze the data using more conventional multivariate statistical approaches, it would be difficult to formalize the hypotheses to be tested, especially considering the variety of strategies that the musicians seem to exhibit. However, in future studies, the two approaches could eventually complement each other: hierarchical clustering can assist in formalizing concrete hypotheses while conventional analyses provide more robust statistical results.

The similarity profiles method (SIMPROF) permits the clustering of statistically similar performances together. It should be noted that other techniques such as bootstrapping [23, 24] may be used as statistical means to define the boundaries of clusters in the mtPs matrix (Figure 2). We leave the exploration of these alternative techniques for a forthcoming publication.

An important parameter in SIMPROF is the significance level (alpha) with which the null hypothesis (that the differences in the mtPs inside a cluster are the result of random combination of the various features) is rejected. The value of alpha, together with the Bonferroni correction for multiple simultaneous tests, determines the level of detail of the final classification, or in other words, the size and scope of the clusters. For instance, if we do not adjust the p-values for multiple tests (Bonferroni correction parameter set to *false* in the `f_disprof_clust` function of the Fathom toolbox [16]), groups B and F are split into three and two sub-clusters respectively. Looking at the mtPs in Figure 2, we see that for group B this is due to the subtle tendency of some performances to play ahead of the beat. In group F, it is due to the relation of the kick and hi-hat strokes. In the more common approach to hierarchical clustering, in which dendrograms are cut horizontally, this level is controlled by the height that a dendrogram is cut.

The way mtPs are clustered together plays central role in the creation of archetypes and therefore in the characterization of the various timing strategies. Archetypes are calculated as averages of the mtPs in each cluster which are further reduced into eight boolean values. For example, if group F was to be split into two sub-clusters, they would not correspond to the same archetype but would form distinct groups. In contrast, the sub-clusters of group B discussed above correspond to the same archetypes.

Exploring and understanding the results of the clustering analysis requires a side-by-side examination of the mtP matrix (Figure 2) and the phylogenetic tree (Figure 3). The two visualizations combined offer an overview of the performances allowing for a closer examination of the finer timing relation details between the different strategies. This simultaneous multilevel view of the recordings enables us to draw conclusions about the timing strategies which would otherwise be obfuscated or oversimplified.

In conclusion, the encoding of the drum performances into boolean features and the hierarchical classification of the derived microtiming profiles effectively cluster the performances into meaningful groups. The further phylogenetic visualization and the symbolic representation of the groups through microtiming archetypes is an efficient way of mapping drummers' main timing strategies, providing an easily interpretable overview of the results. Our analysis simultaneously brings to the surface higher level rhythmic aspects common to the performances as well as the finer details that differentiate them without the one occluding the other.

6. ACKNOWLEDGMENTS

This work was partially supported by the Research Council of Norway through its Centres of Excellence scheme, project number 262762.

7. REFERENCES

[1] M. W. Butterfield, “The Power of Anacrusis: Engendered Feeling in Groove-Based Musics,” *Music Theory Online*, vol. 12, no. 4, pp. 1–17, 2006.

[2] G. S. Câmara, “Swing in Early Funk and Jazz-Funk (1967–1971): Micro-Rhythmic and Macro-Structural Investigations,” University of Oslo, 2016.

[3] G. S. Câmara and A. Danielsen, *Groove*, vol. 1. Oxford University Press, 2018.

[4] G. S. Câmara and A. Danielsen, “Forthcoming.”

[5] K. R. Clarke, P. J. Somerfield, and R. N. Gorley, “Testing of null hypotheses in exploratory community analyses: similarity profiles and biota-environment linkage,” *J. Exp. Mar. Bio. Ecol.*, vol. 366, no. 1–2, pp. 56–69, 2008.

[6] A. Danielsen, *Presence and Pleasure: The Funk Grooves of James Brown and Parliament*. Middletown, CT: Wesleyan University Press, 2006.

[7] A. Danielsen, “Pulse as Dynamic Attending. Analysing Beat Bin Metre in Neo Soul Grooves,” in *The Routledge Companion to Popular Music Analysis: Expanding Approaches*, C. Scotto, K. M. Smith, and J. Brackett, Eds. London: Routledge, 2018, pp. 179–189.

[8] A. Danielsen *et al.*, “Where is the beat in that note? Effects of attack, duration, and frequency on the perceived timing of musical and quasi-musical sounds,” *J. Exp. Psychol. Hum. Percept. Perform.*, vol. 45, no. 3, pp. 402–418, 2019.

[9] A. Danielsen, C. H. Waadeland, H. G. Sundt, and M. A. G. Witek, “Effects of instructed timing and tempo on snare drum sound in drum kit performance,” *J. Acoust. Soc. Am.*, vol. 138, no. 4, pp. 2301–2316, 2015.

[10] A. Friberg and A. Sundström, “Swing Ratios and Ensemble Timing in Jazz Performance: Evidence for a Common Rhythmic Pattern,” *Music Percept.*, vol. 19, no. 3, pp. 333–349, Mar. 2002.

[11] S. Fujii, M. Hirashima, K. Kudo, T. OhtsuKi, Y. NaKamura, and S. Oda, “Synchronization error of drum kit playing with a metronome at different tempi by professional drummers,” *Music Percept.*, vol. 28, no. 1, pp. 491–503, 2011.

[12] W. Goebl and R. Parncutt, “The influence of relative intensity on the perception of onset asynchronies,” in *ICMPC7, July 2002*, 2002, no. July, pp. 1–4.

[13] K. Hellmer and G. Madison, “Quantifying Microtiming Patterning and Variability in Drum Kit Recordings: A Method and Some Data,” *Music Percept. An Interdiscip. J.*, vol. 33, no. 2, pp. 147–162, Dec. 2015.

[14] M. J. Hove, C. Marie, I. C. Bruce, and L. J. Trainor, “Superior time perception for lower musical pitch explains why bass-ranged instruments lay down musical rhythms,” in *Proceedings of the National Academy of Sciences*, 2014, vol. 111, no. 28, pp. 10383–10388.

[15] V. Iyer, “Embodied Mind, Situated Cognition, and Expressive Microtiming in African-American Music,” *Music Percept.*, vol. 19, no. 3, pp. 387–414, Mar. 2002.

[16] D. L. Jones, “Fathom Toolbox for MATLAB: software for multivariate ecological and oceanographic data analysis.” College of Marine Science, University of South Florida, St. Petersburg, FL, USA., 2017.

[17] O. Lartillot, P. Toiviainen, and T. Eerola, “A Matlab Toolbox for Music Information Retrieval,” in *Data Analysis, Machine Learning and Applications. Studies in Classification, Data Analysis, and Knowledge Organization.*, C. Preisach, H. Burkhardt, L. Schmidt-Thieme, and R. Decker, Eds. 2008, pp. 261–268.

[18] P. Legendre and L. Legendre, *Numerical Ecology*, 3rd ed., vol. 24. Elsevier B.V., 2012.

[19] G. Madison, “Experiencing Groove Induced by Music: Consistency and Phenomenology,” *Music Percept.*, vol. 24, no. 2, pp. 201–208, Dec. 2006.

[20] M. Müller and T. Röder, “Motion Templates for Automatic Classification and Retrieval of Motion Capture Data,” in *SCA ’06 Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 2006, pp. 137–146.

[21] B. H. Repp, “Sensorimotor synchronization: A review of the tapping literature,” *Psychon. Bull. Rev.*, vol. 12, no. 6, pp. 969–992, Dec. 2005.

[22] O. Senn, L. Kilchenmann, R. Von Georgi, and C. Bullerjahn, “The effect of expert performance microtiming on listeners’ experience of groove in swing or funk music,” *Front. Psychol.*, vol. 7:1487, no. September, 2016.

[23] R. Suzuki and H. Shimodaira, “Pvclust: An R package for assessing the uncertainty in hierarchical clustering,” *Bioinformatics*, vol. 22, no. 12, pp. 1540–1542, 2006.

[24] R. Suzuki and H. Shimodaira, “An application of multiscale bootstrap resampling to hierarchical clustering of microarray data: How accurate are

these clusters," *15th Annu. Int. Conf. Genome Informatics, Posters Softw. Demonstr.*, 2004.

[25] M. Wojtczak, A. H. Mehta, and A. J. Oxenham, "Rhythm judgments reveal a frequency asymmetry in the perception and neural coding of sound synchrony," *Proc. Natl. Acad. Sci.*, vol. 114, no. 5, pp. 1201–1206, Jan. 2017.