STATISTICAL ANALYSIS OF TEMPO AND RHYTHMIC ELABORATION IN CARNATIC MUSIC

Ajay Srinivasamurthy[1 †], Vignesh Ishwar[⋆] and Xavier Serra[2]

[1] *Amazon.com, Bengaluru, India*
[2] *Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain*
e-mail: ajays.murthy@upf.edu

Statistical analysis of different musical dimensions on a large corpus of music is a valuable tool for computational musicology. Such an analysis can help to supplement and extend observations from careful manual analysis with insights from large corpora of music. We present a statistical analysis of tempo and rhythmic elaboration in Carnatic music using a large corpus of recorded performances. Using a beat-level annotated corpus of Carnatic music, we study the global distribution of tempo across different tāḷas of Carnatic music. Using the median tempo of each piece, we study the local variations in tempo within a music piece. We further analyse the stress patterns within a metrical cycle by averaging spectral features over metrical cycles, to produce cycle-length rhythm patterns. The tempo distributions and the rhythm patterns reveal several insights into their relation to the underlying metrical framework of tāḷa. The observations also provide data based evidence for some aspects of rhythm in Carnatic music that is common knowledge among musicians and musicologists, demonstrating the value of computational methodologies for the analysis of music.

1 Introduction

Current studies in digital humanities largely focus on language and social data, while the study of music corpora has received less attention. Analysis of audio music corpora of music performances can provide us several insights into music and demonstrate the differences and similarities between music theory and practice. The broad goals of analyzing music corpora are to improve our understanding of specific properties of music through empirical observations. Manual analysis of corpora have been an essential part of disciplines such as musicology, but have been limited to a small collection of representative pieces due to the enormous manual effort necessary. However, recent advances in computational methods for analysis enable us to process larger amounts of data more easily. Data-driven analysis of large corpora can provide corpus-level inferences for a musicologist, complementing a manual detailed analysis of small set of representative pieces. Such analyses can also provide us with insights that are difficult to obtain with manual analysis.

Within the CompMusic (`http://compmusic.upf.edu`) project, methods for the analysis of five specific music cultures were developed and Serra [1] introduces the corpora that were compiled for the evaluation of these tools. The criteria for the compilation are motivated by the need to use the corpora for the evaluation of computational analysis methods. In this article, we base our analyses on a corpus that emerged using the

[†] Ajay Srinivasamurthy contributed to this paper during his work period at Music Technology Group, UPF, Barcelona before joining Amazon. [⋆] Vignesh Ishwar is currently a professional Carnatic musician based out of Chennai, India.

guidelines as presented by Serra [1]. We demonstrate how a corpus that originally targeted development in audio processing can be applied to the analysis of structures in music performances, with results relevant to research in the musicologies and digital humanities. The aim of this study is to showcase the presented methods as a potential application of corpus level analysis, while showing their utility for performance analysis and comparative analysis in musicology.

The aim here is not to seek all musicological insights from data, but to illustrate the possibilities of a corpus level analysis data, and how such analysis tools can help aid and advance musicology. These analyses can corroborate several musicological inferences, and can provide additional insights into the differences between musicology, music theory and music practice. At the outset, it is necessary to note that the insights we discuss and conclusions we draw are limited by the available annotated dataset, and hence need further validation. It is however useful to focus on the methodology, which can aid musicologists and engineers to build systems that use these methods for different analyses.

1.1 Motivation and recent work

Carnatic (Karṇāṭaka) music is an art music tradition predominant in the southern part of the Indian subcontinent spanning South India and Sri Lanka. It has a long history of performance and continues to exist and evolve in the current sociocultural contexts. It has a large audience and has attracted a large amount of interest from music scholarship, addressing a variety of questions related to this music culture. The presence of a large dedicated audience and of research literature form a solid basis for studying this music culture from both a musicological and computational perspective. As an orally transmitted and significantly improvised music tradition without concrete music scores, performance analyses on audio recordings is valuable for musicological study of Carnatic music. Recent efforts in curating large amounts of digitally available audio recordings of Carnatic music (CompMusic project, see description by Serra [2]) enables us to perform performance analysis using larger audio corpora. In this work, we focus on an analysis of rhythmic characteristics of Carnatic music.

The recent work in the context of corpus studies can be roughly divided into symbolic and audio based studies. A detailed review of corpus studies in music can be found in [3]. Most corpus studies have focused on western (Eurogenetic) musics and on melodic aspects. Most of the studies doing corpus based analysis in Carnatic music have focused mainly on the aspects of rāga (the melodic framework), e.g. the study of gamaka [4] or intonation [5]. Analysis of rhythm in Carnatic music has been explored from a synthesis perspective [6] or an automatic analysis standpoint addressing Music Information Research (MIR) tasks [7, 8]. Recently, a study focusing on the rhythmic elaboration in Hindustani music showed the value of rhythm analysis of large annotated audio music corpora [3]. We extend the work in [3] to Carnatic music in this paper.

With a sizeable annotated corpus of Carnatic music, we can do corpora level analysis of rhythmic characteristics. We provide a detailed description of the corpus and the tempo distribution in the corpus. Further, we focus on a statistical analysis of rhythm patterns in the Carnatic music corpus. Carnatic music is rhythmically organized within the framework of metrical time cycles called the tāḷa, with the tāḷa cycle being the most important metrical structure. This means that we perform an *intra-cycle* analysis that aims to present typical rhythm patterns as they occur throughout the duration of a tāḷa cycle. We present cycle-length descriptions of rhythmic features that facilitate a visualization of which parts of the cycle are commonly emphasized by the musicians. We start with a brief introduction to rhythm in Carnatic music.

1.2 Rhythm in Carnatic music

Sambamoorthy [9] provides a detailed description of tāḷas in Carnatic music. In Carnatic music, the tāḷa provides a broad structure for repetition of music phrases, motifs and improvisations. It consists of fixed length time cycles called āvartana which can be referred to as the tāḷa cycle. In an āvartana of a tāḷa, phrase refrains, melodic and rhythmic changes occur usually at the beginning of the cycle. An āvartana is divided into basic equidistant time units called akṣaras. The first akṣara pulse of each āvartana is called the sama, which marks the beginning of the cycle (or the end of the previous cycle, due to the cyclic nature of the tāḷa). The sama is often accented, with notable melodic and percussive events. Each tāḷa also has a distinct, possibly non-regular

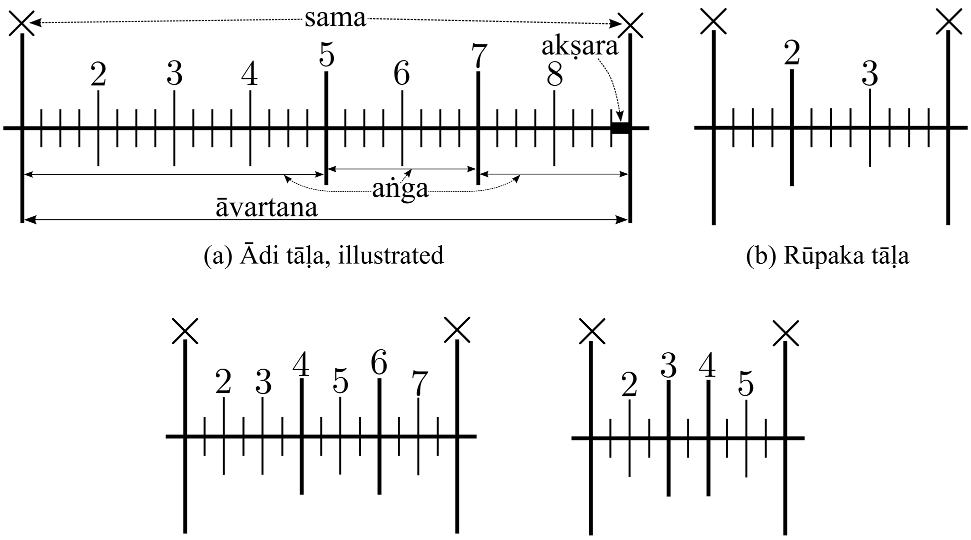

(a) Ādi tāḷa, illustrated

(b) Rūpaka tāḷa

(c) Miśra chāpu tāḷa

(d) Khaṇḍa chāpu tāḷa

Figure 1: An āvartana of four popular Carnatic tāḷas, showing the akṣaras (all time ticks), beats (numbered time ticks), aṅgas (long and bold time ticks) and the sama (×). Ādi tāḷa is also illustrated using the terminology used in this article.

division of the cycle period into sections called the aṅga. The aṅgas serve to indicate the current position in the āvartana and aid the musician to keep track of the movement through the tāḷa cycle. A movement through a tāḷa cycle is explicitly shown by the musician using hand gestures, based on the aṅgas of a tāḷa.

The common definition of an isochronous (equally spaced in time) beat pulsation, as the time instances where a human listener is likely to tap his/her foot to the music [10], is likely to cause problems in Carnatic music. Due the explicit hand gestures, listeners familiar to Carnatic music tend to tap to an non-isochronous sequence of beats in certain tāḷas. Hence we use an adapted definition of a beat for the purpose of a common ground, defined as a uniform pulsation. The akṣaras in an āvartana are grouped into equal length units, which we will refer to as the beats of the tāḷa. The perceptually relevant hand/foot tapping time "beats" are a subset of this uniform beat pulsation. Though there are significant differences in terms of scale and length, as an analogy, the concepts of akṣara, the beat, and the āvartana of Carnatic music bear analogy to the subdivision, beat and the bar metrical levels of Eurogenetic music. Further, aṅga can be interpreted as the possibly unequal length sections of a tāḷa, formed by grouping of beats.

Carnatic music has a sophisticated tāḷa system that incorporates the concepts described above. There are seven basic tāḷas defined with different aṅgas, each with five variants leading to the popular 35 tāḷa system [9]. However, most of these tāḷas are extremely rare in performances with just over a ten tāḷas that can be regularly seen in concerts. A majority of pieces are composed in four popular tāḷas - ādi, rūpaka, miśra chāpu and khaṇḍa chāpu, which are the focus of this paper. The structure of those four popular tāḷas are illustrated in Figure 1. The different concepts related to the tāḷas of Carnatic music are also illustrated in Figure 1(a). The figure shows the akṣaras with time-ticks, beats of the cycle with numbered longer time-ticks, and the sama in the cycle using ×. The aṅga boundaries are highlighted using bold and long time-ticks e.g. ādi tāḷa has 8 beats in a cycle, with 4 akṣaras in each beat leading to 32 akṣaras in a cycle, while rūpaka tāḷa has 12 akṣaras in a cycle, with 4 akṣaras in each of its 3 beats. Some audio examples illustrating these tāḷas and structure of more tāḷas can be seen at `http://compmusic.upf.edu/examples-taala-carnatic`.

Most performances of Carnatic music are accompanied by the percussion instrument mridangam (mr̥daṅgaṁ), a double-sided barrel drum. There could however be other percussion accompaniments such as ghaṭam (the clay pot), khañjira (the Indian tambourine), thevil (a two sided drum) and mōrsiṅg (the Indian jaw harp), which follow the mridangam closely. Since the progression through the tāḷa cycles is explicitly shown through hand gestures, the mridangam is provided with substantial freedom of rhythmic improvisation during the performance. The tāḷa only provides a metrical construct, within which several different rhythmic patterns can be

Table 1: $\mathtt{CMR_f}$ dataset showing the tāḷa structure, total duration and number of annotations. #Sama shows the number of sama annotations and #Ann. shows the number of beat annotations (including samas). $\overline{T_f}$ indicates the median piece length in the dataset (m and s indicate minutes and seconds, respectively)

Tāḷa	#beats per cycle	#Akṣara	#Pieces	Total Duration hours (min)	$\overline{T_f}$	#Ann.	#Sama
Ādi	8	32	50	4.21 (252.78)	4m51s	22793	2882
Rūpaka	3	12	50	4.45 (267.45)	4m37s	22668	7582
Miśra chāpu	7	14	48	5.70 (342.13)	6m35s	54309	7795
Khaṇḍa chāpu	5	10	28	2.24 (134.62)	4m25s	21382	4387
Total			176	16.61 (996.98)	5m4s	121602	22646

Table 2: Tāḷa cycle length indicators for $\mathtt{CMR_f}$ dataset. $\overline{\tau_s}$ and σ_s indicate the mean and standard deviation of the median inter-sama interval of the pieces, respectively. $\overline{\tau_o}$ and σ_o indicate the mean and standard deviation of the median inter-akṣara interval of the pieces, respectively. $[\tau_{s,\text{min}}\ ,\ \tau_{s,\text{max}}]$ indicate the minimum and maximum value of τ_s and hence the range of τ_s in the dataset. All values in the table are in seconds.

Tāḷa	$\overline{\tau_s} \pm \sigma_s$	$\overline{\tau_o} \pm \sigma_o$	$[\tau_{s,\text{min}}\ ,\ \tau_{s,\text{max}}]$
Ādi	5.34 ± 0.723	0.167 ± 0.023	[2.88, 7.07]
Rūpaka	2.13 ± 0.239	0.178 ± 0.020	[1.21, 3.10]
Miśra chāpu	2.67 ± 0.358	0.191 ± 0.026	[1.63, 3.65]
Khaṇḍa chāpu	1.85 ± 0.284	0.185 ± 0.028	[0.91, 2.87]

played and improvised. In this paper, we aim to study these rhythmic patterns and their relationship to the metrical framework of tāḷa. To study those patterns, we built a beat-level annotated corpus of audio music recordings of Carnatic music, that is described in the next section.

2 Carnatic music rhythm dataset

The Carnatic Music Rhythm ($\mathtt{CMR_f}$) dataset (`http://compmusic.upf.edu/carnatic-rhythm-dataset`) is a rhythm annotated test corpus for many automatic rhythm analysis tasks in Carnatic Music [11]. The dataset consists of audio excerpts from the Carnatic research corpus, manually annotated time-aligned markers indicating the progression through the tāḷa cycle, and the associated tāḷa related metadata.

$\mathtt{CMR_f}$ dataset is described in Table 1, showing the four tāḷas and the number of pieces for each tāḷa. The dataset has pieces in four popular tāḷas that encompass a majority of current day Carnatic music performance. The pieces include a mix of vocal and instrumental recordings, recent and old recordings, and span a wide variety of forms. All pieces have a percussion accompaniment, predominantly mridangam. There are also several different pieces by the same artist (or release group), and multiple instances of the same composition rendered by different artists. Each piece is uniquely identified using the MusicBrainz IDentifier (MBID) of the recording. The pieces are mp3 stereo recordings, and sampled at 44.1 kHz. The audio is also available as downmixed mono WAV files for experiments. The audio files are full length pieces or clips extracted from full length pieces. Of the 176 audio files, 120 contain full length pieces. The total duration of audio in the dataset is over 16.6 hours, with 121062 time-aligned beat annotations. The median length of a piece is about 5 minutes in the dataset.

There are several annotations that accompany each excerpt in the dataset. The primary annotations are audio synchronized time-stamps indicating the different metrical positions in the tāḷa cycle - the sama (downbeat) and other beats shown with numerals in Figure 1. The annotations were created using Sonic Visualizer [12] by tapping to music and manually correcting the taps. The annotations have been verified by a professional Car-

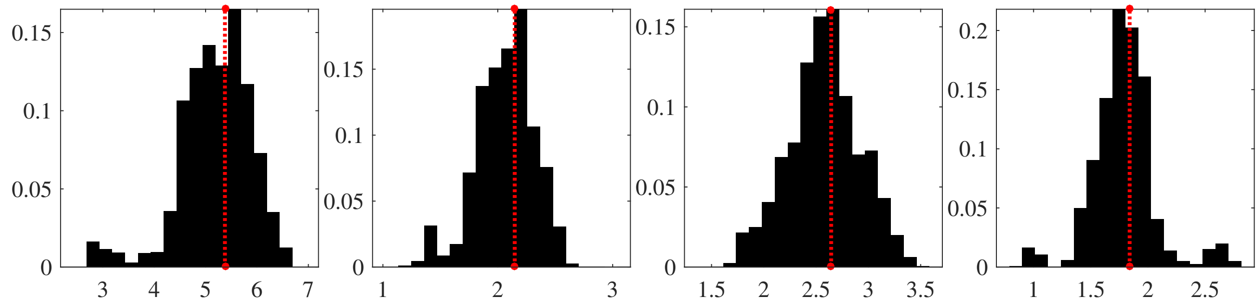

(a) Ādi (b) Rūpaka (c) Miśra chāpu (d) Khaṇḍa chāpu

Figure 2: A histogram of the inter-sama interval τ_s in the $\texttt{CMR}_\mathrm{f}$ dataset for each tāḷa. The ordinate is the fraction of the total count corresponding to the τ_s value shown in abscissa. The median τ_s for each tāḷa is shown as a red dotted line.

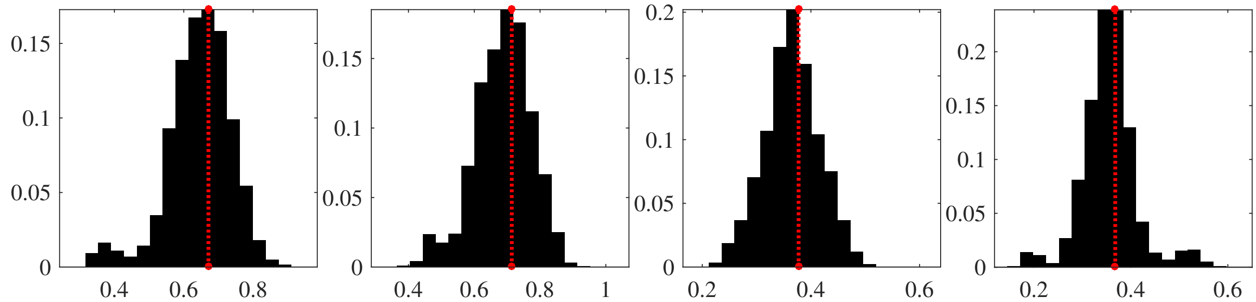

(a) Ādi (b) Rūpaka (c) Miśra chāpu (d) Khaṇḍa chāpu

Figure 3: A histogram of the inter-beat interval τ_b in the $\texttt{CMR}_\mathrm{f}$ dataset for each tāḷa. The ordinate is the fraction of the total count corresponding to the τ_b value shown in abscissa. The median τ_b for each tāḷa is shown as a red dotted line.

natic musician. Each annotation has a time-stamp and an associated numeric label that indicates the position of the beat marker in the tāḷa cycle. In addition, for each excerpt, the tāḷa of the piece is recorded. The possibly time varying tempo of a piece can be obtained using the beat and sama annotations.

The annotations in the $\texttt{CMR}_\mathrm{f}$ dataset are released under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0). The audio in the dataset is copyrighted material sourced from commercially available music releases uniquely identifiable through the MBID and hence is accessible.

3 Tempo distribution in the $\texttt{CMR}_\mathrm{f}$ dataset

Table 2 shows a basic statistical analysis of the tāḷa cycle length indicators in the dataset, which is useful to understand the tempo characteristics and the range of the metrical cycle lengths in the dataset. Ādi tāḷa is the longest tāḷa in the dataset and hence has the highest $\overline{\tau_s}$ among all the tāḷas. Despite no notated tempo, we can see from the values of the median inter-akṣara interval, $\overline{\tau_o}$ and its standard deviation that the tempo in Carnatic music does not vary much across the tāḷas. The range of $\overline{\tau_s}$ values show that a wide range of cycle durations that are present in Carnatic music pieces. The shortest cycle in the dataset is less than second long, while the longest cycle is over 7 seconds long.

The tempo values are not notated in Carnatic music, and the pieces are not played to a metronome. Hence, in addition to the median values tabulated in Table 2 we present further analysis of the inter-sama interval (τ_s) and inter-beat interval (τ_b) for each tāḷa over the whole $\texttt{CMR}_\mathrm{f}$ dataset. A histogram of τ_s and τ_b for each tāḷa is shown in Figure 2 and Figure 3 respectively. This shows the distribution of cycle lengths in the dataset over the whole range of τ_s for each tāḷa, around the median value. Despite the large range of τ_s values, the distribution

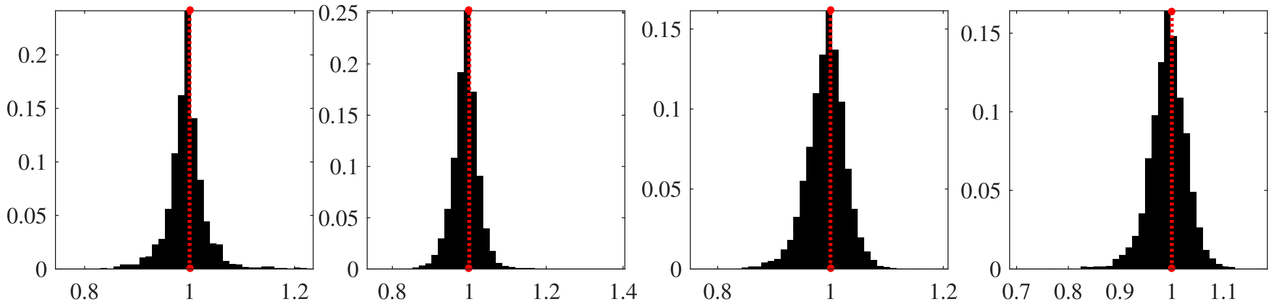

(a) Ādi (b) Rūpaka (c) Miśra chāpu (d) Khaṇḍa chāpu

Figure 4: A histogram of the median normalized inter-sama interval τ_s in the $\mathtt{CMR_f}$ dataset for each tāḷa. The ordinate is the fraction of the total count corresponding to the normalized τ_s value shown in abscissa.

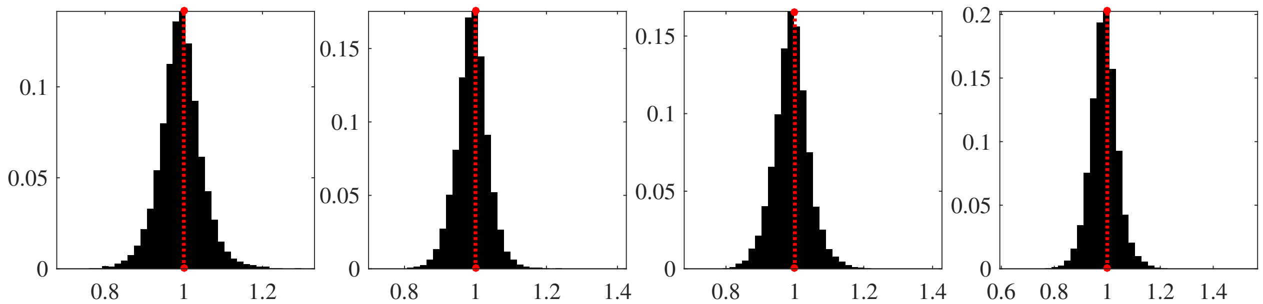

(a) Ādi (b) Rūpaka (c) Miśra chāpu (d) Khaṇḍa chāpu

Figure 5: A histogram of the median normalized inter-beat interval τ_b in the $\mathtt{CMR_f}$ dataset for each tāḷa. The ordinate is the fraction of the total count corresponding to the normalized τ_b value shown in abscissa.

in Figure 2 and Figure 3 show that the tempo often is limited to a small range of values. Though the musicians are free to choose any tempo, we empirically observe that they tend to choose a narrow range of tempo.

To illustrate and measure the time varying tempo of music pieces in Carnatic music, we normalize all the τ_s and τ_b values in a piece by the median value of the piece to obtain median normalized τ_s and τ_b values, a histogram of which is shown for $\mathtt{CMR_f}$ dataset in Figure 4 and Figure 5, respectively. These histograms are centered around 1 since they are normalized by the median, and the spread of these histograms around the value of 1 is a measure of deviation of tempo from the median value. From the figures, it is clear that the tempo is time varying but with less than about 20% maximum deviation from the median tempo of the piece for all tāḷas. This is an indicator of tempo variations in a music piece and of expressive timing that is often a part of Carnatic music performances. The tempo deviations are seen to be similar across all tāḷas.

4 Cycle level rhythm analysis of the $\mathtt{CMR_f}$ dataset

The āvartana (tāḷa cycle) is the most relevant metrical level in the tāḷa, and the level around which the whole Carnatic music performance is organized, and hence an analysis of rhythm at cycle-level is relevant. Cycle-level rhythm patterns aim to depict typical stress patterns that occur in the various parts of the tāḷa cycle and can be analyzed in relation with the underlying tāḷa progression. These rhythm patterns are computed automatically and are strongly related to the strokes of the percussion instrument. We average the spectral flux features organized into cycle length sequences to compute a canonical cycle length rhythm pattern for each tāḷa. The process is explained in detail in the next section.

4.1 Computation of canonical cycle length rhythm patterns

The rhythm patterns are computed using a feature proposed in [13] for the scope of detecting musical onsets in audio recordings. Since this feature is derived from the time-derivative of the Short-Time Fourier Transform (STFT) magnitude (i.e. the spectrum), it can be referred to as a *spectral flux* feature. Such spectral flux features are generally motivated by the fact that the onsets of musical events, such as percussion strokes or a singer intoning a new note, are accompanied by energy increases in certain frequency regions in the spectrum of the signal. The term Spectral Flux expresses this idea of quantifying energy fluctuations in the spectral domain. Furthermore, Spectral Flux features have been successfully applied for the task of automatic meter analysis from audio recordings using rhythm patterns in e.g., the work by [14] and [15] and hence is a suitable feature for analysis of rhythm patterns.

To compute the features, the STFT of the audio signal with a window size of 46.4 ms (2047 samples of audio at a sampling rate of 44.1 kHz), DFT size of 2048 and hop size of 20 ms is computed from audio. The successive difference between frames of the logarithm of the filter bank energies in 82 different bands is then computed. Since the bass onsets have significant information about the rhythmic patterns, the features are computed in two frequency bands (Low: $\leq$ 250 Hz, High: $>$ 250 Hz) to additionally consider the bass onsets. The process of computing the spectral flux feature is illustrated in Figure 6.

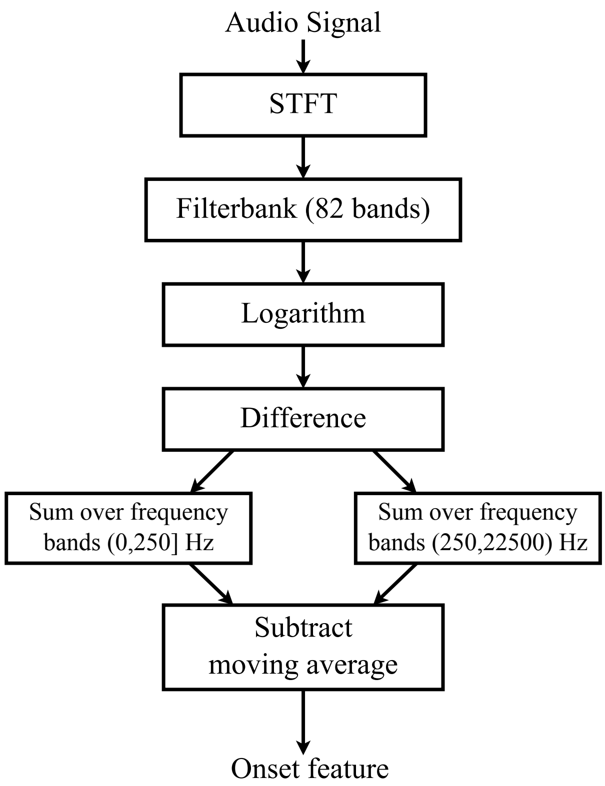

Figure 6: Computation of the spectral flux onset feature in two frequency bands, from [16].

Using beat and downbeat annotated training data, the spectral flux features from all music pieces in a specific tāḷa are then grouped into cycle length sequences, and interpolated to equal lengths using a fine grid. A mean of all such cycle length sequence instances for a specific tāḷa is computed in both the frequency bands and used as a canonical rhythmic pattern illustrated here.

At the outset, it is necessary to note here that the patterns played in a tāḷa cycle are better described using timing, energy and timbre descriptors. The rhythm patterns generated here using the spectral flux feature and can only explain timing and energy accents. A minor effect of timbre can be seen in these rhythm patterns, but are predominantly affected by the other two characteristics. These patterns are averaged over the whole dataset for a tāḷa, and hence cannot capture specific nuances of individual pieces, but only can give a corpus-level perspective. The patterns here are indicative of the surface rhythm present in the audio recordings, and hence completely reflect the underlying canonical metrical structures.

4.2 Observations from cycle length rhythm patterns

Figures 7-10 show the canonical cycle length patterns over all the pieces in the dataset for each tāḷa, computed using the spectral flux feature in two different frequency bands as outlined in Section 4.1. In each figure, the bottom pane corresponds to the low frequency band and the top pane corresponds to the high frequency band. The abscissa is the beat number within the cycle (dotted lines), with 1 indicating the sama (marked with a red line). The start of each aṅga is indicated with beat numbers at the top of each pane (sama shown as $\times$). The patterns in each figure pane is normalized so that maximum value is 1, to comment on relative onset strengths at different metrical positions of the cycle.

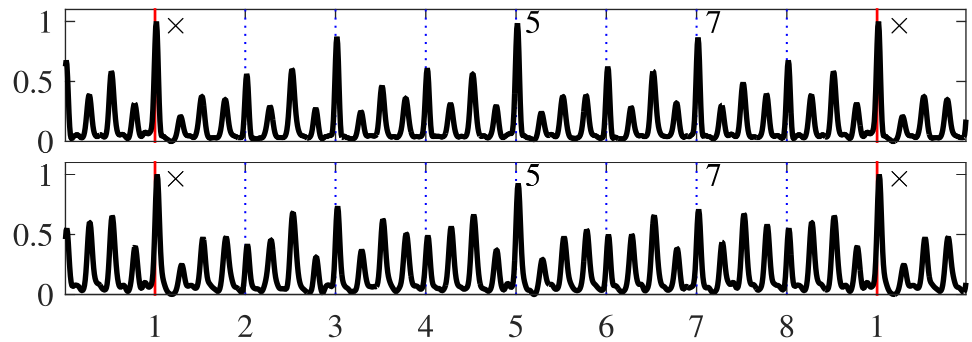

Figure 7: Cycle length rhythmic patterns learned from $\mathtt{CMR_f}$ dataset for ādi tāḷa. In each of the following Figures 7-10, the patterns are computed from spectral flux feature and averaged over all the pieces in the dataset. The bottom/top pane corresponds to the low/high frequency bands, respectively. The abscissa is the beat number within the cycle (dotted lines), with 1 indicating the sama (marked with a red line). The start of each aṅga is indicated with beat numbers at the top of each pane (sama shown as $\times$). The plot shows the cycle extended by a beat at the beginning and end to illustrate the cyclic nature of the tāḷa.

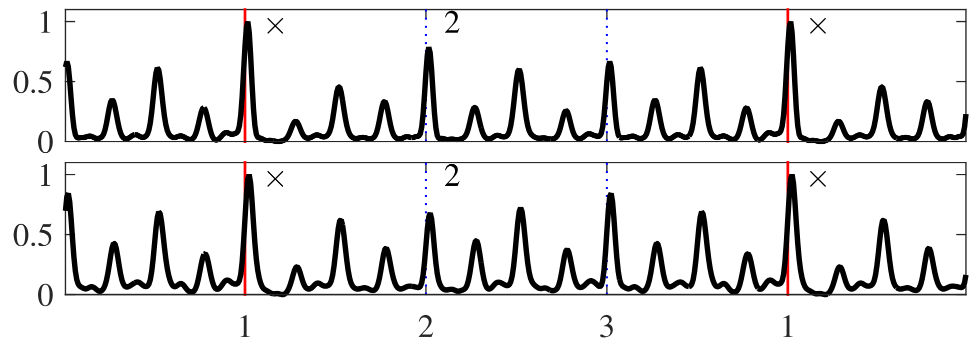

Figure 8: Cycle length rhythmic patterns learned from $\mathtt{CMR_f}$ dataset for rūpaka tāḷa.

The rhythm patterns are indicative of the energies of mridangam strokes played in the cycle. In the figures, the bottom pane that shows the low frequency band has content from the left bass drum while the top pane has content predominantly from the right pitched drum (and additionally from the lead melody). Hence, for the purpose of this discussion, we use the terms left and right accents to refer to the accents in rhythm patterns shown on the bottom and top pane, respectively. The left and right accents provide interesting insights into the patterns played within a tāḷa cycle. Finally, it is important to note that the patterns illustrated here are average patterns that occur and hence do not tell us much about the various individual patterns that might occur in specific points in particular recordings.

We list down and discuss some salient qualitative observations from figures Figures 7-10. Overall, we see stronger accents on the akṣaras, with sama having the strongest accent in most cases. We can clearly see the accents organized in three different strengths, reflecting the metrical levels of the aṅga, the beat and the akṣara. The two akṣara long beats in miśra chāpu and khaṇḍa chāpu tāḷas, and the four akṣara long beats in ādi and rūpaka tāḷas can be additionally seen. The patterns played in Carnatic music are quite diverse, and no obvious representative tāḷa pattern can be inferred, apart from the varied accents at three metrical levels. The tāḷas are metrical structures that allow many different patterns to be played, and not a specific rhythm. It is further seen that the first akṣara after sama has softer accents. Fewer strokes are played after the sama, to emphasize that the sama has just passed and a new cycle has begun. It might also perhaps indicate some form of recovery time after the intense stroke-playing towards the end of the cycle.

We now discuss some tāḷa specific observations. Figure 7 shows the rhythm patterns for ādi tāḷa. We see that a three level hierarchy of aṅga, beats and akṣaras is well demarcated. The akṣara at half cycle (beat 5) has an accent as strong as the sama. The odd beats (marked 1, 3, 5, 7) have stronger right accents. The left accents are distributed through the cycle, with strong accents at half cycle. Figure 8 shows the rhythm patterns for rūpaka tāḷa. Apart from the three level hierarchy of accents that is quite apparent, the half beat accent between the beats 2 and 3 are strong - indicating the often played 6+6 akṣara grouping structure of rūpaka, with a ternary meter.

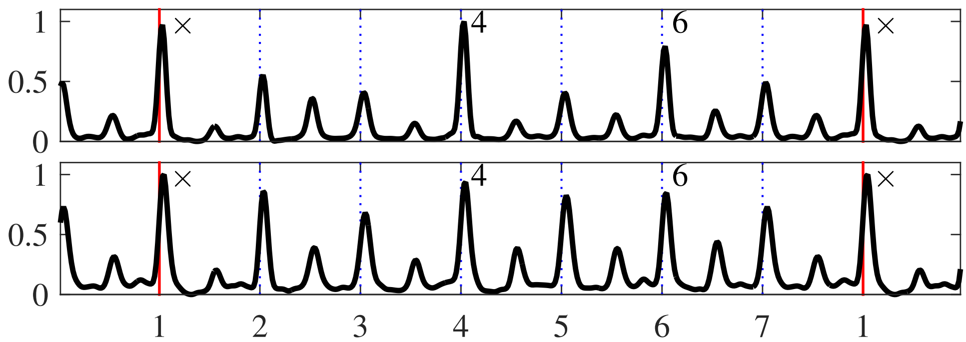

Figure 9: Cycle length rhythmic patterns learned from $\mathtt{CMR_f}$ dataset for miśra chāpu tāḷa.

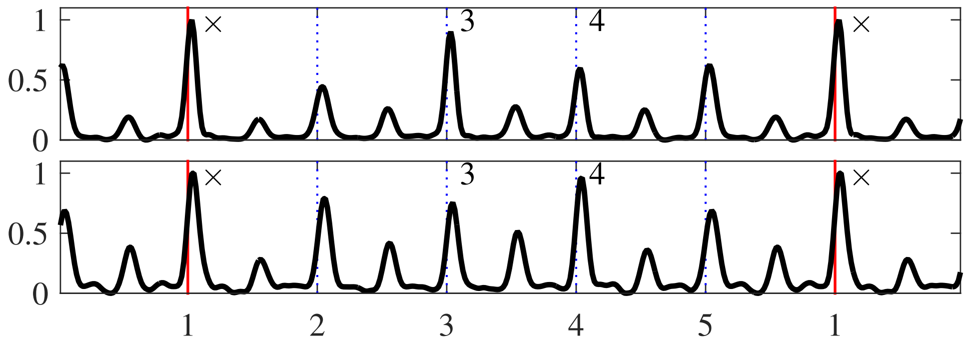

Figure 10: Cycle length rhythmic patterns learned from $\mathtt{CMR_f}$ dataset for khaṇḍa chāpu tāḷa.

Figure 9 shows the rhythm patterns for miśra chāpu tāḷa. We see that the aṅga boundaries have strong left and right accents showing their use as anchor points to indicate the progression through the cycle. Though defined with a 3+2+2 akṣara grouping structure, a 1+2+2+2 structure is often seen in miśra chāpu tāḷa, which can be observed here, based on the strong left accent on beat 2. An additional strong left accent on beat 6 shows that it is also used as an anchor. The rhythm patterns of khaṇḍa chāpu tāḷa shown in Figure 10 have a strong left accent on beat 4, which is used as an anchor within the cycle. A stronger right accent on beat 3 shows the progression through the unequal aṅgas. The 2+1+2 akṣara grouping structure of khaṇḍa chāpu is often played out as 3+2 or 2+3, showing strong accents on beats 3 and 4.

The overall and tāḷa specific observations presented above from rhythm patterns have interesting musicological significance. Most of these observations are common knowledge among musicians and musicologists and the data-driven analysis presented here corroborates it. A professional Carnatic musician has validated these observations, but there are many more insights that might be derived from these patterns with a further in-depth analysis of cycle length rhythm patterns, leading to valid musicological conclusions.

5 Conclusions

We presented a corpus-level statistical analysis of tempo and rhythmic elaboration in Carnatic music. Starting with a beat-level annotated dataset of audio music recordings of Carnatic music spanning over four common tāḷas, we studied the inter-sama and inter-beat intervals to analyze the global distribution and local variations of tempo in music pieces. We observed that the median tempo used by musicians is similar across different tāḷa, while upto around 20% variation in inter-sama and inter-beat interval is seen in the dataset from the corresponding median values. We used the spectral flux features in two freuqncy bands to computed the canonical cycle length rhythm patterns. These rhythm patterns provided insights into stress and accents at different parts of a tāḷa cycle, in relation to the underlying metrical structure of the tāḷa. These observations corroborate common knowledge among musicians and musicologists about some aspects of rhythmic elaboration in Carnatic music, showing the utility of the large corpus of data and analyses tools described in the paper to derive musicological insights. While the paper only presented some insights, the methodology allows for deeper analysis to derive further musically meaningful insights on the complex aspects of rhythm in Carnatic music.

Acknowledgments

This work is partly funded by the European Research Council under the European Union's Seventh Framework Program, as part of the CompMusic project (ERC grant agreement 267583). Parts of the paper are from the PhD dissertation of Ajay Srinivasamurthy [7].

REFERENCES

[1] X. Serra, "Creating Research Corpora for the Computational Study of Music: the case of the CompMusic Project," in *Proceedings of the 53rd AES International Conference on Semantic Audio*, London, 2014.

[2] X. Serra, "A multicultural approach in Music Information Research," in *Proceedings of the 12th International Society for Music Information Retrieval Conference (ISMIR 2011)*, Miami, USA, 2011, pp. 151–156.

[3] A. Srinivasamurthy, A. Holzapfel, K. K. Ganguli, and X. Serra, "Aspects of Tempo and Rhythmic Elaboration in Hindustani Music: A Corpus Study," *Frontiers in Digital Humanities*, vol. 4, no. 20, 2017.

[4] V. S. Viraraghavan, R. Aravind, and H. A. Murthy, "A Statistical Analysis of Gamakas in Carnatic Music." in *Proceedings of the 18th International Society for Music Information Retrieval Conference (ISMIR 2017)*, Suzhou, China, Oct. 2017, pp. 243–249.

[5] G. K. Koduri, J. Serrà, and X. Serra, "Characterization of intonation in carnatic music by parametrizing pitch histograms," in *International Society for Music Information Retrieval Conference (ISMIR 2012)*, Porto, Oct. 2012, pp. 199–204.

[6] K. Trochidis, C. Guedes, A. Holzapfel, A. Akshay, and A. Klaric, "Analysis-by-synthesis of rhythm in south indian art percussion performances by means of statistical analysis," in *Proceedings of the International Conference on Analytical Approaches to World Music (AAWM)*, New York, USA, Jun. 2016.

[7] A. Srinivasamurthy, "A Data-driven Bayesian Approach to Automatic Rhythm Analysis of Indian Art Music," Doctoral dissertation, Universitat Pompeu Fabra, Nov. 2016.

[8] A. Srinivasamurthy, A. Holzapfel, and X. Serra, "Informed Automatic Meter Analysis of Music Recordings," in *Proceedings of the 18th International Society for Music Information Retrieval Conference (ISMIR 2017)*, Suzhou, China, Oct. 2017, pp. 679–685.

[9] P. Sambamoorthy, *South Indian Music Vol. I-VI*. The Indian Music Publishing House, 1998.

[10] A. Holzapfel, M. Davies, J. R. Zapata, J. L. Oliveira, and F. Gouyon, "Selective Sampling for Beat Tracking Evaluation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 9, pp. 2539–2548, Nov. 2012.

[11] A. Srinivasamurthy and X. Serra, "A Supervised Approach to Hierarchical Metrical Cycle Tracking from Audio Music Recordings," in *Proceedings of the 39th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2014)*, Florence, Italy, May 2014, pp. 5237–5241.

[12] C. Cannam, C. Landone, and M. Sandler, "Sonic Visualiser: An Open Source Application for Viewing, Analysing, and Annotating Music Audio Files," in *Proceedings of the ACM Multimedia 2010 International Conference*, Florence, Italy, Oct. 2010, pp. 1467–1468.

[13] S. Böck, F. Krebs, and M. Schedl, "Evaluating the Online Capabilities of Onset Detection Methods," in *Proceedings of the 13th International Society for Music Information Retrieval Conference (ISMIR 2012)*, Porto, Portugal, 2012, pp. 49–54.

[14] F. Krebs, A. Holzapfel, A. T. Cemgil, and G. Widmer, "Inferring Metrical Structure in Music Using Particle Filters," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 5, pp. 817–827, May 2015.

[15] A. Srinivasamurthy, A. Holzapfel, A. T. Cemgil, and X. Serra, "Particle Filters for Efficient Meter Tracking with Dynamic Bayesian Networks," in *Proceedings of the 16th International Society for Music Information Retrieval Conference (ISMIR 2015)*, Malaga, Spain, Oct. 2015, pp. 197–203.

[16] A. Holzapfel, F. Krebs, and A. Srinivasamurthy, "Tracking the "odd": Meter inference in a culturally diverse music corpus," in *Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR 2014)*, Taipei, Taiwan, 2014, pp. 425–430.