

DISCOVERING RĀGA MOTIFS BY CHARACTERIZING COMMUNITIES IN NETWORKS OF MELODIC PATTERNS

Sankalp Gulati[†], Joan Serrà^{*}, Vignesh Ishwar[†], Xavier Serra[†]

[†]Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

^{*}Telefonica Research, Barcelona, Spain

ABSTRACT

Rāga motifs are the main building blocks of the melodic structures in Indian art music. Therefore, the discovery and characterization of such motifs is fundamental for the computational analysis of this music. We propose an approach for discovering rāga motifs from audio music collections. First, we extract melodic patterns from a collection of 44 hours of audio comprising 160 recordings belonging to 10 rāgas. Next, we characterize these patterns by performing a network analysis, detecting non-overlapping communities, and exploiting the topological properties of the network to determine a similarity threshold. With that, we select a number of motif candidates that are representative of a rāga, the rāga motifs. For a formal evaluation we perform listening tests with 10 professional musicians. The results indicate that, on an average, the selected melodic phrases correspond to rāga motifs with 85% positive ratings. This opens up the possibilities for many musically-meaningful computational tasks in Indian art music, including human-interpretable rāga recognition, semantic-based music discovery, or pedagogical tools.

Index Terms— Rāga motif, melodic phrases, melodic similarity, Indian art music, Carnatic music

1. INTRODUCTION

Indian art music (IAM) comprises two music traditions, Hindustani music [1] and Carnatic music [2]. Both these music traditions are heterophonic in nature, with their melodies being based on the rāga framework [3]. Rāgas are characterized by a set of svaras (roughly speaking, notes), ārōha-avrōha (the ascending and descending melodic progression) and, most importantly, by a set of characteristic melodic phrases referred to as rāga motifs. Their functional role is to act as building blocks to construct melodies, and as a base for the artist to express his/her creativity through improvisation within the rāga grammar. Rāga motifs capture the essence of a rāga, and are the prominent cues for its identification [4]. Discovery of rāga motifs can pave the way for efficient rāga-based music retrieval from large audio archives, semantically meaningful music discovery and navigation, and several applications around music pedagogy.

A rāga motif, in textbook notation, is represented as a sequence of svaras. However, there is no explicit set of rules for rāga motifs, and its realization in a musical performance is subject to the artist’s interpretation and creativity. The high degree of pitch and timing variation across the occurrences of such motifs makes the task of pattern discovery very challenging in IAM. The lack of melody segmentation models and the large duration of audio recordings further complicate this task. As a result, existing research on this topic within

IAM is primarily focused on computing robust melodic representations and similarity measures [5–11]. Descriptive music scores and a reliable melody transcription system are practically non-existent for IAM. Therefore, the majority of these approaches work on a continuous pitch-based melody representation. Though recent proposals exploiting culture specificity for improving similarity look promising [8, 12], their overall accuracy is still far from perfect. This poses a challenge to the pattern characterization process, as the output of the pattern discovery contains many musically-meaningless matches.

Another challenge in discovering rāga motifs is the presence of other categories of repeating melodic patterns. Though a rāga motif recurs in a rāga rendition, not all repeating patterns are rāga motifs. We can broadly classify such patterns other than the rāga motifs into two categories: (a) patterns that occur across different rāgas, which are not specific to a rāga or a composition, here called gamaka motifs¹, and (b) patterns that are specific to a composition within a rāga but do not occur in any another composition, here called composition-specific motifs. Neither of these two type of melodic patterns can be used to categorize or identify a particular rāga. Hence, we here focus on the characterization of melodic patterns in order to distinguish rāga motifs from other types of motifs.

There are a number of existing approaches for the discovery and characterization of melodic motifs in diverse contexts, and for different music traditions [13–19]. A majority of these approaches rely on a symbolic music representation. This could be attributed to the challenges involved in the extraction of a reliable melodic representation from raw audio recordings [19], and to the easy availability of descriptive music scores in these music traditions. In IAM, there is no standardized symbolic representation of melodies. Furthermore, the concepts used by musicians to describe melodic motifs are difficult to capture in a symbolic representation. Due to the fundamental differences between the western-classical music based representation, used in most studies, and the specificities of IAM, we need an approach not based on symbolic representation of the melodies.

We present here a novel approach to characterize the melodic patterns, that are discovered in audio music collections of IAM, in order to identify rāga motifs. For extracting melodic patterns, we make use of an existing unsupervised approach (Section 2.1). We then propose to perform a network analysis to cluster these patterns into communities (Section 2.2). In this process, we also present an approach to determine a similarity threshold, which is yet another challenge in such tasks. We then empirically devise a measure to rank the detected communities according to rāga representativeness (Section 2.2.3). Finally, we validate the top-ranked communities by performing listening tests with professional musicians (Section 4). The results are very encouraging, with 85% positive ratings.

This work is partly supported by the European Research Council under the European Unions Seventh Framework Program, as part of the CompMusic project (ERC grant agreement 267583).

¹For example Kampita, a specific type of an oscillatory melodic movement on a svar [4].

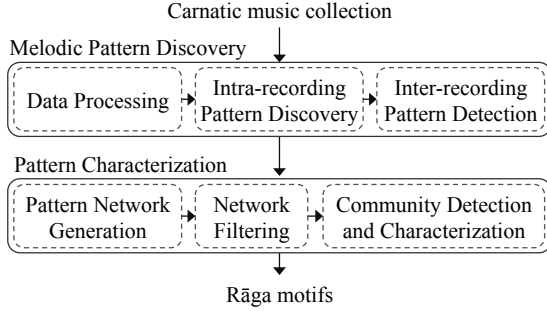


Fig. 1. Block diagram of the proposed approach.

2. METHOD

The block diagram for the proposed approach is shown in Figure 1. There are two main blocks: (a) melodic pattern discovery and (b) pattern characterization. Note that, we do not refer to discovered melodic patterns (Section 2.1) as motifs, since they only represent repeating melodic fragments, with no specific musical identity at this stage. In the pattern characterization step (Section 2.2) we cluster these patterns and identify the ones that represent rāga motifs.

2.1. Melodic Pattern Discovery

As mentioned, discovering melodic patterns in audio recordings of IAM is a challenging task. To get a reliable input for our approach, we employ the IAM pattern discovery system we presented in [10]. This is one of the few unsupervised systems, we are aware of, that can discover meaningful melodic patterns in large-scale collections of IAM. We use the same parameter settings and the implementation of the method as reported in the paper. This pattern discovery method consists of three main processing blocks (Figure 1). The method uses tonic normalized predominant pitch as the melody representation, and considers all possible 2 second melodic segments as the pattern candidates in the data processing block. This duration is chosen based on the average human annotated phrase duration reported in recent studies [12]. In the intra-recording pattern discovery block, the method discovers melodically similar seed pattern pairs within each audio recording. Subsequently, in the inter-recording pattern detection block, every seed pattern is then used to perform a search across all the recordings in the dataset to obtain its nearest neighbors. This approach uses dynamic time warping based melodic similarity, and employs several lower-bounding techniques to reduce the computational complexity of the task. For a detailed explanation of this approach we refer to [10].

2.2. Melodic Pattern Characterization

As explained, a discovered melodic pattern can correspond to either a rāga motif, a composition-specific motif or to a gamaka motif (Section 1). The aim of the pattern characterization block is to cluster these discovered patterns and to characterize the obtained clusters in order to identify the ones that represent different rāga motifs. For this, we propose to perform a network analysis as described below.

2.2.1. Pattern Network Generation

We start by building an undirected weighted network using the discovered patterns from the previous step. The patterns are considered as the nodes of the network and the edge between any two patterns (i ,

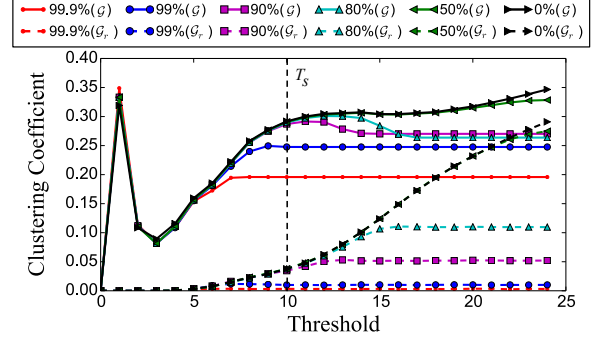


Fig. 2. Evolution of the clustering coefficient of \mathcal{G} and \mathcal{G}_r over different thresholds and for different statistical confidence values used for disparity filtering (legend).

j) is weighted based on the distance D_{ij} between the patterns. Noticeably, D_{ij} is computed using the same distance measure as used in the pattern discovery block of Section 2.1. The weight of the edge W_{ij} between the nodes i and j is given by $W_{ij} = e^{-D_{ij}/D}$, where, D is the mean of D_{ij} over every combination of i and j .

2.2.2. Network Filtering

The main objective of this processing block is to filter the network to retain only the musically meaningful connections between the nodes. Since the edge weights between the pairs of melodically similar and dissimilar nodes may vary by orders of magnitude, we first consider to exploit this heterogeneity to extract the network’s backbone. We therefore need to apply disparity filtering [20] to preserve only the edges that represent statistically significant deviations with respect to a null model of edge weight assignment for every node. The only parameter used in the disparity filtering is the statistical confidence value. We iterate over 5 different confidence values $\{99.99, 99, 90, 80, 50\}$. However, as we will show, the application of disparity filtering is found to be quite irrelevant for the present case.

We next proceed to filter edges in the network based on a melodic similarity threshold T_S . We propose to estimate T_S based on the topological properties of the network. For this, we analyze the evolution of the clustering coefficient of both the obtained network \mathcal{G} and the corresponding randomized network \mathcal{G}_r over a range of similarity thresholds. Clustering coefficient measures the extent to which the nodes in a network tend to cluster together [21]. The randomized network \mathcal{G}_r is obtained by swapping the edges between randomly selected node pairs such that the degree of each node is preserved [22]. This way, \mathcal{G}_r can be considered as the maximally random network with that particular degree distribution. In Figure 2, we show the evolution of the clustering coefficient of \mathcal{G} and \mathcal{G}_r over different similarity thresholds (indicated by exponentially spaced bins). In addition, we can also see the clustering coefficient curves for different statistical confidence values used for disparity filtering. The evolution of the clustering coefficients is used for obtaining a similarity threshold as explained below.

We hypothesize that the more musically meaningful T_S is, the higher is the difference between the clustering coefficients of \mathcal{G} and \mathcal{G}_r . We therefore select $T_S = 10$. Note that even though the similarity threshold corresponding to $T_S = 1$ results in a higher value of the clustering coefficient, we reject it because the filtered network consists only of a small number of nodes. These nodes correspond to near-exact pattern repetitions discovered within the same recording [10]. Such patterns typically represent composition-specific mo-

tifs, and hence are irrelevant in our context. In Figure 2, we also observe that the disparity filtering using a confidence value higher than 80% significantly lowers the clustering coefficient, which can be attributed to the removal of musically meaningful edges in the network. On the other hand, given $T_S = 10$, the disparity filtering with a confidence value lower than 80% does not significantly affect the clustering coefficient. We can thus conclude that, in the given scenario, disparity filtering does not bring in any clear advantage. Finally, after applying T_S , we transform \mathcal{G} to an unweighted network.

2.2.3. Community Detection and Characterization

We next take the unweighted undirected network that results from the previous step, and perform a non-overlapping community detection using the method proposed in [23]. This method is based on modularity optimization and is parameter-free from the point of view of the user. It has been extensively used in various applications [24] and can deal with very large networks [23]. We use the implementation available in networkX [25], a Python language package for exploration and analysis of networks and network algorithms. Using this method for our entire dataset, we obtain around 1800 communities of melodic patterns.

A community C_q is comprised of N nodes, and the node count over different rāgas is given by the ordered list $\alpha_q = (\alpha_{q,1}, \alpha_{q,2}, \dots, \alpha_{q,L})$ such that $\alpha_{q,i} \geq \alpha_{q,j}, \forall i < j$, where each element in α_q denotes the number of nodes in a particular rāga and L is the total number of unique rāgas comprising the community. Similarly, the node count over the audio recordings is given by the ordered list $\beta_q = (\beta_{q,1}, \beta_{q,2}, \dots, \beta_{q,K})$ such that $\beta_{q,l} \geq \beta_{q,m}, \forall l < m$, where each element in β_q denotes the number of nodes belonging a particular audio recording and K is the total number of recordings comprising the community. For both these cases, $\sum_{i=1}^L \alpha_{q,i} = \sum_{l=1}^K \beta_{q,l} = N$.

We now proceed to characterize the detected communities in order to identify the ones that represent rāga motifs. For that we first categorize a community C_q as belonging to the rāga R_q corresponding to the maximum number of nodes $\alpha_{q,1}$ in that community. Subsequently, for each rāga, we rank all the communities belonging to that rāga. To do so we empirically devise a goodness measure γ , which denotes the likelihood that a community C_q represents a rāga motif. We propose to use

$$\gamma = N\rho^4\lambda, \quad (1)$$

where ρ is an estimate of the likelihood of rāga R_q in C_q ,

$$\rho = \frac{\alpha_{q,1}}{N}, \quad (2)$$

and λ indicates how uniformly the nodes of the community are distributed over audio recordings,

$$\lambda = \frac{\sum_{l=1}^K l \cdot \beta_{q,l}}{N}. \quad (3)$$

Higher λ implies a more uniform distribution. Since a community that represents a rāga motif is expected to contain nodes from a single rāga (high value of ρ) and the nodes belong to many different recordings (high value of λ), the goodness measure γ is high for such a community. In general we prefer large communities, but, to avoid detecting large communities (high value of N) corresponding to gamaka motifs (low value of ρ) we use a fourth power on ρ . Composition-specific motifs are expected to have a low λ , as they are not repeated across multiple recordings.

Rāga	Dur	#Com	#Art	Svaras
Hamsadhvani	2.46	12	14	$s r_2 g_3 p n_3$
Kāmavardhini	3.94	13	16	$s r_1 g_3 m_2 p d_1 n_3$
Darbār	2.59	8	13	$s r_2 g_2 m_1 p d_2 n_2$
Kalyāṇi	6.94	9	16	$s r_2 g_3 m_2 p d_2 n_3$
Kāmbōji	6.91	12	13	$s r_2 g_3 m_1 p d_2 n_2 n_3$
Bēgaḍa	3.41	9	16	$s r_2 g_3 m_1 p d_2 n_2 n_3$
Kāpi	2.24	12	16	$s r_2 g_2 m_1 p d_2 n_2 n_3$
Bhairavi	5.33	7	16	$s r_2 g_2 m_1 p d_2 d_3 n_2$
Behāg	1.51	12	16	$s r_2 g_3 m_1 m_2 p d_2 n_2 n_3$
Tōḍi	8.75	12	16	$s r_1 g_2 m_1 p d_1 n_2$
Total	44.08	106	57	-

Table 1. Details of the dataset in terms of the duration (Dur) in hours, number of unique compositions (#Com), unique lead artists (#Art), and the svaras for each rāga. Here s, r, g, m, p, d, n denote the 7 svaras in IAM and the subscript indicates the variant of the svar for a particular rāga (cf. [2]).

3. EVALUATION

3.1. Music Collection

The music collection used in this study is compiled as a part of the CompMusic project [26, 27]. The collection comprises 44 hours of polyphonic audio music recordings of Carnatic music across 10 different rāgas. For each rāga we select 16 music pieces, which amounts to a total of 160 recordings. There are 139 vocal music recordings and 21 instrumental recordings comprising violin, vīṇa and bamboo flute. In Table 1, we summarize the relevant details of the dataset. We see that it is diverse in terms of the number of unique compositions and number of lead artists. Furthermore, it includes different forms of compositions (kīrtana, varnam and viruttam) and recordings containing varied improvised sections such as ālāpāna, nereval and kalpanā-svaras. The chosen rāgas contain diverse set of svaras (notes) both in terms of the number of svaras and their pitch classes (svarasthānās). From Table 1, we also notice that several rāgas such as Kalyāṇi, Kāmbōji and Bēgaḍa have a large fraction of svaras in common. We refer to them as allied rāgas. This further increases the complexity of the task at hand, since the discrimination between the phrases of allied rāgas may be based on subtle melodic nuances.

3.2. Setup and Evaluation Measures

Given the unsupervised nature of this study, we perform a listening test to formally evaluate the extent to which the selected melodic phrases correspond to rāga motifs. For each of the 10 rāgas in the dataset, we select the top 10 communities based on the goodness measure γ (Eq. 1). From each of these communities, we select their representative melodic phrase based on the betweenness centrality of the nodes [21], i.e., the node with the highest betweenness centrality is considered as the representative melodic phrase of that community. In case of a tie, we select the one with the highest node degree. Finally, we arrive at a set of 100 melodic phrases, which are then used to perform the listening test. These audio examples are also made available online².

For the listening test we select 10 professional Carnatic musicians with over 15 years of formal music training. Each musician is presented with the audio fragments corresponding to the selected

²<http://compmusic.upf.edu/node/277>

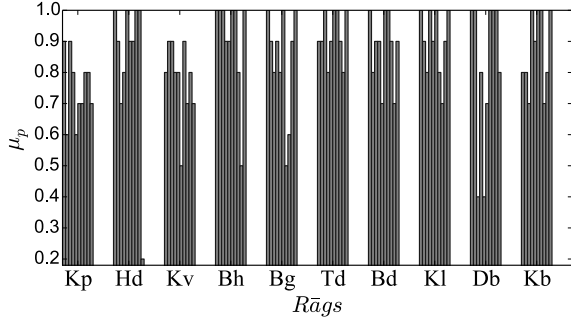


Fig. 3. Mean musician rating per melodic phrase for each rāga: Kāpi (Kp), Hamsadhvani (Hd), Kāmavardhini (Kv), Bhairavi (Bh), Behāg (Bg), Tōḍi (Td), Bēgādā (Bd), Kalyāṇī (Kl), Darbār (Db), Kāmbōji (Kb).

melodic phrases in a random order. They are also presented with the rāgas corresponding to the melodic phrases. The musicians are asked to rate each melodic phrase based on whether it is a characteristic phrase of that rāga. We use binary ratings (‘Yes’ or ‘No’).

The audio fragments were segmented with a one second buffer on either side of the phrase to offer some context and reduce the effect of abrupt boundaries. In order to quantify the musicians’ assessment, we use mean ratings for each phrase p , μ_p , considering ‘Yes’ as 1 and ‘No’ as 0. For analyzing the ratings per rāga, we study the mean and standard deviation of all μ_p for phrases in every rāga, which we denote by μ_r and σ_r , respectively.

4. RESULTS AND DISCUSSION

We first analyze the musicians’ ratings at the level of melodic phrases. In Figure 3, we show μ_p for the 100 selected melodic phrases, where the grouping is based on their corresponding rāgas. We find that the mean and the standard deviation of μ_p for the melodic phrases is 0.85 and 0.16, respectively. For a better understanding of μ_p across phrases and the overall musicians’ agreement, we show the histogram of μ_p in Figure 4. We see that 33 melodic phrases are rated as rāga motifs by all 10 musicians and 25 phrases are rated as rāga motifs by 9 out of 10 musicians. Similarly, the musicians’ agreement can be inferred for the rest of the phrases from this histogram. We observe that 91% of the phrases are always marked as rāga motifs by at least 7 out of 10 musicians.

We now proceed to analyze the results for different rāgas. In Table 2, we summarize mean μ_r and standard deviation σ_r of μ_p for each rāga. We observe that there is a considerable amount of variation in μ_r across the rāgas. It ranges from 0.75 for rāga Kāpi to 0.92 in the case of rāga Tōḍi. An interesting observation here is that the phrase-based rāgas³ are the top performing rāgas with the exception of rāga Darbār. From Table 2 and Table 1, we notice a strong correlation between μ_r and the total duration of the audio recordings across rāgas. This suggests that longer music pieces are likely to facilitate the discovery of rāga motifs owing to more number of occurrences of such melodic phrases.

We now examine the melodic phrases with low scores. An investigation of 9 out of 100 phrases that obtain $\mu_p \leq 0.6$ reveals that many of these phrases are composition-specific phrases that do not characterize the rāga. The method wrongly identifies them as rāga motifs because their associated communities have a high γ score owing to a high λ value. This can be attributed to the fact that these

³Rāgas whose identity is derived based on phraseology than svaras [4].

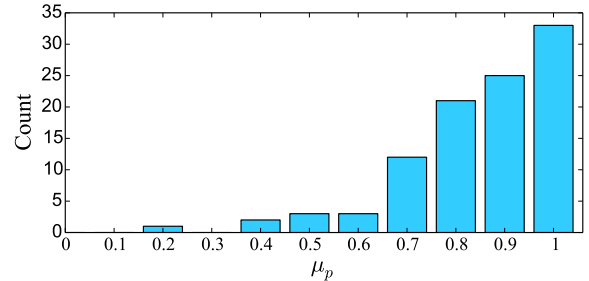


Fig. 4. Histogram of μ_p for all 100 melodic phrases.

Rāga	μ_r	σ_r	Rāga	μ_r	σ_r
Hamsadhvani	0.84	0.23	Bēgādā	0.88	0.11
Kāmavardhini	0.78	0.17	Kāpi	0.75	0.10
Darbār	0.81	0.23	Bhairavi	0.91	0.15
Kalyāṇī	0.90	0.10	Behāg	0.84	0.16
Kāmbōji	0.87	0.12	Tōḍi	0.92	0.07

Table 2. Mean μ_r and standard deviation σ_r of μ_p for each rāga. Rāgas with $\mu_r \geq 0.85$ are highlighted.

phrases are discovered from multiple recordings, since their corresponding compositions have several instances in the dataset. This suggests that, possibly, the goodness measure γ can be made more robust to such cases by computing λ using the distribution of nodes over unique compositions rather than over audio recordings.

The results show that the proposed method successfully discovers rāga motifs with a high accuracy. We see that, even for the allied rāgas present in the dataset such as Kāmbōji and Bēgādā (Section 3.1), the method is able to discover distinct characteristic rāga motifs. As mentioned, allied rāgas are challenging because they have a substantial overlap in the set of svaras that they comprise (see also Table 1). Finally, on a more informal side, it is worth mentioning that musicians were impressed when, after the listening test, they came to know that the melodic phrases were discovered by a machine following an unsupervised approach.

5. CONCLUSION AND FUTURE WORK

We presented a novel unsupervised approach to discover rāga motifs from polyphonic audio music collections of IAM and, specifically, to distinguish them from gamaka and composition-specific motifs. We first extracted melodic patterns from audio recordings using an existing unsupervised approach. We then employed a network analysis and non-overlapping community detection algorithm to cluster melodic patterns. Using the topological properties of the network, we determined a musically meaningful similarity threshold. In addition, we devised a goodness measure for characterizing the detected communities. We evaluated our method using a sizable and representative music collection. A listening test with 10 professional Carnatic musicians shows that the proposed method successfully discovers rāga motifs with accuracy, even in the presence of allied rāgas in the dataset. This suggests that the functional roles of different melodic phrases in IAM can be effectively exploited to identify them in an unsupervised manner. This, to the best of our knowledge, has been done for the first time in this study. In the future, we plan to extend this work to identify other melodic phrase categories such as composition-specific motifs. Furthermore, we plan to quantitatively evaluate the discovered melodic phrases by using them in tasks such as rāga recognition and composition identification.

6. REFERENCES

- [1] A. Danielou, *The ragas of Northern Indian music*, Munshiram Manoharlal Publishers, New Delhi, 2010.
- [2] T. Viswanathan and M. H. Allen, *Music in South India*, Oxford University Press, 2004.
- [3] S. Bagchee, *Nād understanding raga music*, Business Publications Inc, 1998.
- [4] T. M. Krishna and V. Ishwar, “Karnāṭic music: Svāra, gamaka, motif and rāga identity,” in *Proc. of the 2nd CompMusic Workshop*, 2012, pp. 12–18.
- [5] J. C. Ross, T. P. Vinutha, and P. Rao, “Detecting melodic motifs from audio for Hindustani classical music,” in *Proc. of Int. Conf. on Music Information Retrieval (ISMIR)*, 2012, pp. 193–198.
- [6] J. C. Ross and P. Rao, “Detection of raga-characteristic phrases from Hindustani classical music audio,” in *Proc. of 2nd Comp-Music Workshop*, 2012, pp. 133–138.
- [7] P. Rao, J. C. Ross, K. K. Ganguli, V. Pandit, V. Ishwar, A. Bellur, and H. A. Murthy, “Classification of melodic motifs in raga music with time-series matching,” *Journal of New Music Research*, vol. 43, no. 1, pp. 115–131, Jan. 2014.
- [8] V. Ishwar, S. Dutta, A. Bellur, and H. Murthy, “Motif spotting in an Alapana in Carnatic music,” in *Proc. of Int. Conf. on Music Information Retrieval (ISMIR)*, 2013, pp. 499–504.
- [9] S. Dutta and H. A. Murthy, “Discovering typical motifs of a raga from one-liners of songs in Carnatic music,” in *Int. Society for Music Information Retrieval (ISMIR)*, Taipei, Taiwan, 2014, pp. 397–402.
- [10] S. Gulati, J. Serrà, V. Ishwar, and X. Serra, “Mining melodic patterns in large audio collections of indian art music,” in *Int. Conf. on Signal Image Technology & Internet Based Systems - MIRA*, Marrakesh, Morocco, 2014, pp. 264–271.
- [11] S. Gulati, J. Serrà, and X. Serra, “An evaluation of methodologies for melodic similarity in audio recordings of indian art music,” in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, 2015, pp. 678–682.
- [12] S. Gulati, J. Serrà, and X. Serra, “Improving melodic similarity in indian art music using culture-specific melodic characteristics,” in *Int. Soc. for Music Information Retrieval Conf. (ISMIR)*, Málaga, Spain, 2015, pp. 680–686.
- [13] D. Conklin, “Discovery of distinctive patterns in music,” *Intelligent Data Analysis*, vol. 14, pp. 547–554, 2010.
- [14] D. Conklin and C. Anagnostopoulou, “Comparative Pattern Analysis of Cretan Folk Songs,” *Journal of New Music Research*, vol. 40, no. 2, pp. 119–125, 2010.
- [15] J. L. Hsu, C. C. Liu, and A. L. P. Chen, “Discovering Nontrivial Repeating Patterns in Music Data,” *IEEE Transactions on Multimedia*, vol. 3, no. 3, pp. 311–325, 2001.
- [16] Z. Juhász, “Motive identification in 22 folksong corpora using dynamic time warping and self organizing maps,” in *Int. Society for Music Information Retrieval*, 2009, pp. 171–176.
- [17] I. Knopke and F. Jürgensen, “A System for Identifying Common Melodic Phrases in the Masses of Palestrina,” *Journal of New Music Research*, vol. 38, no. 2, pp. 171–181, 2009.
- [18] O. Lartillot and M. Ayari, “Motivic pattern extraction in music, and application to the study of Tunisian modal music,” *South African Computer Journal*, vol. 36, pp. 16–28, 2006.
- [19] T. Collins, S. Böck, F. Krebs, and G. Widmer, “Bridging the audio-symbolic gap: The discovery of repeated note content directly from polyphonic music audio,” in *Audio Engineering Society Conference: 53rd International Conference: Semantic Audio*. Audio Engineering Society, 2014.
- [20] M. A. Serrano, M. Boguñá, and A. Vespignani, “Extracting the multiscale backbone of complex weighted networks,” in *Proc. of the National Academy of Sciences of the USA*, vol. 106, no. 16, pp. 6483–6488, 2009.
- [21] M. EJ Newman, “The structure and function of complex networks,” *Society for Industrial and Applied Mathematics (SIAM) review*, vol. 45, no. 2, pp. 167–256, 2003.
- [22] S. Maslov and K. Sneppen, “Specificity and stability in topology of protein networks,” *Science*, vol. 296, no. 5569, pp. 910–913, 2002.
- [23] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre, “Fast unfolding of communities in large networks,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, no. 10, pp. P10008, 2008.
- [24] S Fortunato, “Community detection in graphs,” *Physics Reports*, vol. 486, no. 3, pp. 75–174, 2010.
- [25] A. A. Hagberg, D. A. Schult, and P. J. Swart, “Exploring network structure, dynamics, and function using NetworkX,” in *Proceedings of the 7th Python in Science Conference (SciPy2008)*, Pasadena, CA USA, Aug. 2008, pp. 11–15.
- [26] X. Serra, “A multicultural approach to music information research,” in *Proc. of Int. Conf. on Music Information Retrieval (ISMIR)*, 2011, pp. 151–156.
- [27] A. Srinivasamurthy, G. K. Koduri, S. Gulati, V. Ishwar, and X. Serra, “Corpora for music information research in indian art music,” in *Int. Computer Music Conf./Sound and Music Computing Conf.*, Athens, Greece, 2014, pp. 1029–1036.