Raga Space Visualization: Analyzing Melodic Structures in Carnatic and Hindustani Music

Soham Korade
IIIT Hyderabad
soham.korade@students.iiit.ac.in

Suswara Pochampally
Independent Researcher

Saroja TK
IIIT Hyderabad
saroja.tk@iiit.ac.in

Abstract

The concept of raga in Indian classical music serves as a complex, multifaceted melodic entity that can be approached through various perspectives. Compositions within a raga act as foundational structures, serving as the bedrock for improvisations. Analyzing their textual notations is easier and more objective in comparison with analyzing audio samples. A significant amount of musical insights can be derived from the discrete swara sequences alone.

This paper aims to construct an intuitive visualization of raga space[1], using swara sequences from raga compositions. Notations from public sources are normalized, and their TF-IDF features are projected into a low-dimensional space. This approach allows for qualitative analysis of both Carnatic and Hindustani ragas, mapping them to known raga theory.

1 Introduction

Indian classical music, comprising of Carnatic and Hindustani systems, centers on the concept of *raga*. Ragas are melodic frameworks for improvisation, each with unique characteristics. Compositions objectively represent a raga, encapsulating its grammar and essential phrases in their notations.

Analyzing ragas from notations offers several advantages. It eliminates issues related to the limited availability of high-quality, authentic recordings and avoids the subjectivity inherent in various renditions and interpretations of ragas.

However, this approach has some limitations. Notations do not capture the subtleties of gamakas and other embellishments. Carnatic ragas with multiple varieties of the same notes need to be dismissed due to the inherent ambiguity in the notation.

The core idea of this paper is the visualization of ragas as entities in low-dimensional space, with their sequences as discrete points. Unlike previous works that analyze a small number of ragas, this study examines a wide variety (341) from both systems of Indian classical music, aiming to automate the process with minimal manual work.

This visualization-based comparative analysis helps understand melodic characteristics and similarities/differences within and across both systems, offering a novel perspective on raga structures.

2 Previous Work

(Ross et al., 2017) also employ notations and not audio to identify the similarities between ragas. They use LSTM (Long Short-Term Memory) networks to extract the feature embeddings. The features are *learned*. In contrast, we use a deterministic approach and obtain the features directly from the sequences, making the features interpretable.

(Ganguli et al., 2016) adopts a data-driven methodology to validate existing music concepts pertaining to ragas using audio recordings from Hindustani music concerts.

(Sahasrabuddhe and Upadhye, 1992) modeled a raga as a finite state automaton based on the swara patterns followed in it. (Pandey et al., 2003) extended this idea of swara sequence working with Hidden Markov Models on swara sequences extracted using a heuristics driven note-segmentation technique. They employed a novel pakad[2]-matching algorithm that improved the HMM based results.

3 Data Collection

3.1 Datasets

The dataset for this study was scraped and compiled from several reliable websites related to Hindustani and Carnatic music.

---

[1]Code and dataset are available at:
https://drive.google.com/file/d/
1CfbhpN-HufPERvKo8dezXeNslSO4COpg/view?usp=
sharing

[2]a set of phrases that captures the essence of a raga

Popular	S	r	R	g	G	m	M	P	d	D	n	N
Ours	s	R	r	G	g	m	M	p	D	d	N	n
Semitones	0	1	2	3	4	5	6	7	8	9	10	11
Hindustani	S	R	R	G	G	M	M^l	P	D	D	N	N
Carnatic 1	S	R1 / G1	R2 / G2	R3 / G3		M1	M2	P	D1	D2 / N1	D3 / N2	N3
Carnatic 2	S	R / G	R / G	R / G		M	M	P	D	D / N	D / N	N

Table 1: Swara notation format

Figure 1: Sequence length distribution

Vishwamohini (Sawant, 2024) advertises itself as an online notations library of raga-based and *tala*(rhythm)-based compositions and is free and open to all for contribution and use. The website hosts more than 1000 compositions, 362 are raga-based, 215 of which are usable.

Tanarang (Ringe and Ringe, 2024) hosts comprehensive details of 120 ragas, notation was extracted from *aroh-avroh*, *pakad* and description.

Carnatic Notations (Jeyaraman, 2024) is a blog containing 970 compositions, 844 of which can be used for this research. We want the entire composition to be set to a single raga, so 11 *raga malikas*[3] are excluded.

3.2 Cleaning and Preprocessing

In the Carnatic system of notation, the variety of swara (e.g. R has R1, R2, R3) is not notated, but understood from the context (**Carnatic 2**, see Table 1). Usually, the raga name is mentioned and the variety is deduced from the *aroh-avroh* of the raga, which is notated in **Carnatic 1** format. For ragas which employ a single specific variety for all its notes, we can directly map the generic swara to the specific swara (**Case A**). There are other cases where in the composition, there can be *anya* swaras[4] (**Case B**) or the raga itself employs two varieties of a swara (e.g., D1 and D2 for D) (**Case C**). For cases **B** and **C**, as the swaras need to be manually mapped, we drop the compositions from the dataset.

After data collection, notation formats specific to each source were then converted into our own normalized format. Subsequently, the data underwent cleaning where obviously incorrect compositions and compositions of insufficient length were removed by an expert.

Duration of each note was made constant. All note embellishments were removed. Only the sequence of notes, along the octave markers was pre-

---

[3] meaning "garland of ragas": compositions wherein various segments are set to different ragas
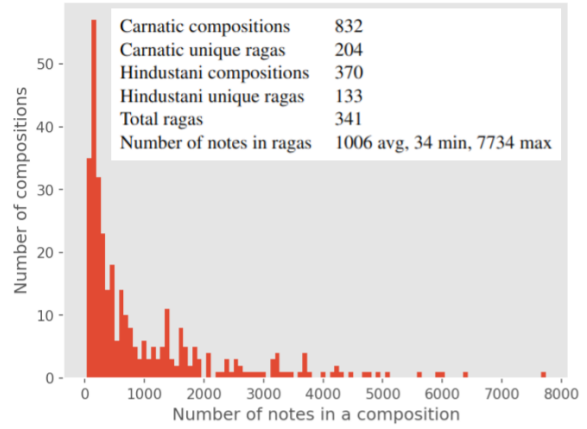
[4] swaras not present in the *aroh-avroh*

served. All the swaras used in the compositions fall into three octaves thus, giving us a vocabulary of $12 \times 3 = 36$ unique swaras.

The raga names were normalized (bhoop, bhupali, bhoopali → bhoopali). The names for carnatic ragas were prefixed with "C_" to avoid name clashes. The compositions were grouped by raga and merged. See Figure 1.

Format Name	Notations

Vishwamohini	`[P1|-|G1|-|D1|P1|`
	`G1|-|S2|S2|D1|`
	`P1|G1|R1|S1|R1]`
	`[notations]`

Tanarang	`P,G,DPG,S'S'DPGRSR`

CarnaticNotations	`mOhanam - S R2 G3 P D2`
	`S' - S' D2 P G3 R2 S`
	`P,G,DPG,S'S'DPGRSR`

Normalized	`pgdpg's'sdpgrsr`

Table 2: Example notation in all formats

4 Feature Extraction

Term Frequency-Inverse Document Frequency (TF-IDF) (Sammut and Webb, 2010) is a numerical statistic that reflects the importance of a term in a document relative to a collection of documents, often used in information retrieval and text mining. In the context of raga sequences, TF-IDF can be employed to extract characteristic "words". Here, we consider swaras as letters, and the ragas as sentences with inherent meaning. Similar ideas were seen in (Garcia-Valencia, 2020).

Given the absence of explicit word boundaries in raga sequences, we enumerate all n-grams. This results in a vocabulary V of words, where V =

{All possible combinations of the 36 unique swaras having length n}.

Here, w_i is a specific n-gram from the vocabulary V and r_j is a raga from the set of all ragas R. The TF-IDF score for a word w_i in a raga r_j is calculated as:

$$\text{TF-IDF}(w_i, r_j) = \text{TF}(w_i, r_j) \times \text{IDF}(w_i) \quad (1)$$

$$\text{where, } \text{TF}(w_i, r_j) = \frac{\text{frequency of } w_i \text{ in } r_j}{\text{\#words in } r_j}$$

$$\text{IDF}(w_i) = \log \left(\frac{\text{total number of ragas}}{\text{\#ragas having } w_i} \right)$$

We use TF-IDF because it is simple and efficient, and has an ability to rank words in a way that is easy to interpret.

5 Methodology

Feature extraction was conducted using the `TfidfVectorizer` class from the `sklearn.feature_extraction.text` module (Pedregosa et al., 2011). A custom tokenizer specific to our notation was built, and the parameter `lowercase` was set to `false` to preserve the integrity of our notations, as case changes would change the swara variety. We found that the parameter `3-5` for generating n-grams was best suited for the experiment, as *pakad*s have words of similar size.

The resulting document-term matrix represents each raga's word composition as a feature vector. To validate the selected features, we rank each raga's word list by TF-IDF score and filter out words not present in the raga sequence. The words hence obtained act as good representation of the ragas, much like those in the *sancharam* or *chalan* as given in (Krishnamacharyulu, 2003; Garg, 2021, 2022). A sample is given in Table 3. The success of this approach also highlights the notion of music as a language.

5.1 Dimensionality Reduction using t-SNE

t-SNE (t-Distributed Stochastic Neighbor Embedding) (van der Maaten and Hinton, 2008; mat, 2018) is a statistical method for visualizing high-dimensional data in a low-dimensional space. It is particularly effective at capturing the non-linear relationships between data points.

In the study, the TF-IDF features are visualized on a 2D plane using t-SNE. This visualization allows for a comprehensive understanding of the relationships within the ragas.

Raga	#words	Top Words
C_shankarabharanam	3601	pmg mgr dns gmp
C_kalyani	2858	dpM pMg ndp snd
C_mayamalavagowla	2323	Rgm mgR sRg
C_thodi	2224	mGR DNs NsR
yaman	2151	Mgr rgM ndp nrg

Table 3: Top 5 Ragas with Most Important Words

5.2 Visualization

The 'Interactive Graph' tool[5] allows entering any new sequence and locating the resulting point in raga space. For example, if we enter "'sDNpGmrs", we see that the new sequence is plotted near Adana and Darbari Kanada.

Using the cosine distance threshold option in the interface one can highlight the closest ragas in the raga space which are within the given distance. On clicking on a raga point, the tool shows more information about the raga. There is an option to play the notation available in the database in the twin tool 'Composer'[6].

6 Results and Discussion

6.1 Analysis of Ragas in the Raga Space

With the assistance of a trained musician specialized in both Carnatic and Hindustani music, we observed several insightful patterns and relationships in the plot. This subsection analyzes insights from Figure 2.

Parallels across systems (*durga-C_shuddha saveri*, *hansadhwani-C_hamsadhwani*, *bhoopali-C_mohanam*), allied ragas (*C_darbar-C_nayaki*, *bhoopali-deshkar*) and similar ragas (*desh-tilak kamod*) are close in the raga space. See (Garg, 2021, 2022; Sambamurthy, 1953)

The Venn diagrams, based on the sets of words, show the relative similarities of the ragas and account for their proximity in the graph. See Figure 3.

6.2 Validation with previous work

The neighbors identified through our method, specifically of marwa, puriya, sohni and yaman, shuddha kalyan and shankara align very closely with known theory. See Figure 4. These neighbors correspond with those obtained by MDS visual-
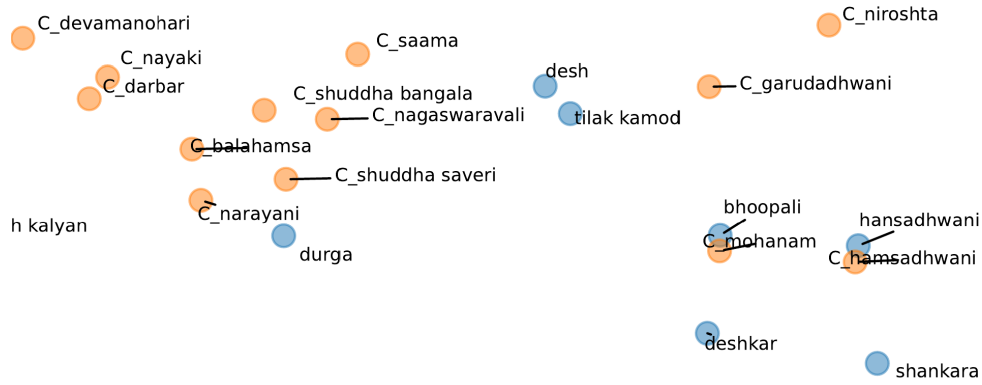
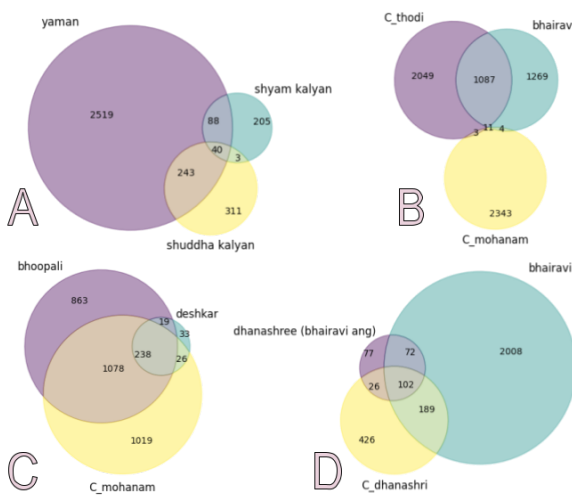Figure 2: A zoomed-in portion of the Raga Space

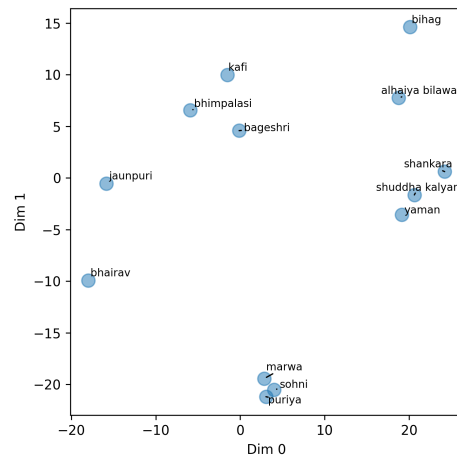Figure 3: Venn diagrams for sets of raga "words"

Figure 4: Comparison of our plot with previous work

ization of bi-LSTM note embeddings similarities mentioned in (Ross et al., 2017).

7   Future Work

A prospective direction for future research lies in the development of a new classification system for ragas grounded in the raga space. This system would integrate both categorization paradigms of Carnatic and Hindustani music, namely the Melakarta and Thaat systems, into a cohesive framework. Such an approach would involve the construction of an empirically-based, mathematically-derived classification system.

Given that this study relies solely on musical notation in melodic systems, it can be extended to incorporate other musical systems, such as the Arabic Maqam system.

8   Conclusion

The visualizations obtained solely from the notations sourced from websites present compelling results that are consistent with established raga theory. Given the dynamic nature of the evolution of new ragas, future ragas naturally integrate into the raga space, serving as reference points for comparison with existing ragas.

The 'Interactive Graph' can be used in music education to give valuable insights on ragas in an engaging manner.

References

2018. *Interactive Exploration of Musical Space with Parametric t-SNE*. Zenodo.

Kaustuv Kanti Ganguli, Sankalp Gulati, Xavier Serra, and Preeti Rao. 2016. Data-driven exploration of melodic structure in hindustani music. In *Devaney J, Mandel MI, Turnbull D, Tzanetakis G, editors. ISMIR 2016. Proceedings of the 17th International Society*

for Music Information Retrieval Conference; 2016 Aug 7-11; New York City (NY).[Canada]: ISMIR; 2016. p. 605-11. International Society for Music Information Retrieval (ISMIR).

Sebastian Garcia-Valencia. 2020. Embeddings as representation for symbolic music. *arXiv preprint arXiv:2005.09406*.

L. Garg. 2021. *Raag Visharad - 1*, 7th edition. Sangeet karyalaya, Hathras, India.

L. Garg. 2022. *Raag Visharad - 2*, 6th edition. Sangeet karyalaya, Hathras, India.

Balaji Jeyaraman. 2024. Carnatic music notations.

NCH Krishnamacharyulu. 2003. *Sangitha Ragadarshini*. Rohini Publications, Rajahmundry, India.

Gaurav Pandey, Chaitanya Mishra, and Paul Ipe. 2003. Tansen: A system for automatic raga identification. In *IICAI*, pages 1350–1363.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Prakash Vishwanath Ringe and Vishwajeet Vishwanath Ringe. 2024. Tanarang.

Joe Cheri Ross, Abhijit Mishra, Kaustuv Kanti Ganguli, Pushpak Bhattacharyya, and Preeti Rao. 2017. Identifying raga similarity through embeddings learned from compositions' notation. In *ISMIR*, pages 515–522.

H Sahasrabuddhe and R Upadhye. 1992. On the computational model of raag music of india. In *Workshop on AI and music: European conference on AI*.

P. Sambamurthy. 1953. *South Indian Music*. The Indian Music Publishing House, Madras, India.

Claude Sammut and Geoffrey I. Webb, editors. 2010. *TF–IDF*, pages 986–987. Springer US, Boston, MA.

Shivraj Sawant. 2024. Vishwamohini.

Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605.