A Generative Model of Ragas Using Directed Graphs

Authors: Bhushan Agarwal
Submitted: 21. July 2025
Published: 25. August 2025
Volume: 12
Issue: 4
Languages: English
Keywords: Melodic Graph; Raga Representation; Carnatic Music; Computational Musicology; Music Generation.
Categories: Performing Arts, Music
DOI: 10.17160/josha.12.4.1064

Abstract:

This paper introduces melodic graphs—directed graphs with vertices representing individual notes and paths representing sequences of notes—to model note transitions in a context-specific and phrase-driven way. The generative capacity of melodic graphs is demonstrated here through the Carnatic raga Haṃsadhvani. A melodic graph was created by representing 16 short phrases in the raga as a connected graph with directed edges. A sample of 10 phrases generated using the melodic graph and 10 random sequences were evaluated by 8 musicians trained in Carnatic music. The phrases were rated on a scale of 1-10, with 1 representing no fit, and 10 representing a good fit of the phrase within the raga. The generated phrases received a significantly higher average rating than the random sequences (average rating of graph-generated sequences: 8.2/10; vs. random sequences: 3.6/10; $p << 0.01$). Thus, the melodic graph models Haṃsadhvani effectively. Future work with other ragas and using more sophisticated melodic graphs is needed to fully explore their applications in music education, raga representation and modelling, and computer-assisted music generation.

JOSHA Journal of Science, Humanities and Arts

August 2025 Volume 12, Issue 4

A Generative Model of Ragas Using Directed Graphs

Bhushan Kumar Agarwal
bhushan141@gmail.com
Independent Scholar, Hyderabad, India

Abstract

This paper introduces melodic graphs—directed graphs with vertices representing individual notes and paths representing sequences of notes—to model note transitions in a context-specific and phrase-driven way. The generative capacity of melodic graphs is demonstrated here through the Carnatic raga Haṃsadhvani. A melodic graph was created by representing 16 short phrases in the raga as a connected graph with directed edges. A sample of 10 phrases generated using the melodic graph and 10 random sequences were evaluated by 8 musicians trained in Carnatic music. The phrases were rated on a scale of 1-10, with 1 representing no fit, and 10 representing a good fit of the phrase within the raga. The generated phrases received a significantly higher average rating than the random sequences (average rating of graph-generated sequences: 8.2/10; vs. random sequences: 3.6/10; $p << 0.01$). Thus, the melodic graph models Haṃsadhvani effectively. Future work with other ragas and using more sophisticated melodic graphs is needed to fully explore their applications in music education, raga representation and modelling, and computer-assisted music generation.

Keywords: Melodic Graph; Raga Representation; Carnatic Music; Computational Musicology; Music Generation.

Introduction

Ragas are commonly delineated using a phrase-based approach through svar-vistār, a series of phrases that show note-patterns that are characteristic of and permissible in a raga [1]. Key melodic phrases that characterise a raga form the building blocks for composition and improvisation [2] in both Hindustani and Carnatic music. No scientific model has yet been formulated to codify the rules governing phrasing in ragas [3]. Note sequences have been modelled elsewhere using finite-state models and Markov models (e.g., [4] and [5]). However, because these state machines account for individual note transitions alone and not for the history of transitions within a phrase, they may not be useful in representing phrases while preserving their local musical context.

Melodic graphs—directed graphs with vertices representing individual notes (svaras) and paths representing phrases in a raga—are proposed in this work as a novel method to model note transitions in a context-specific and phrase-driven way. The vertices representing notes are connected with directional edges representing note transitions, and the path between any 2 vertices represents a melody or a melodic phrase.

Methods

Sixteen short phrases/segments were manually curated from a Carnatic raga, Haṃsadhvani (basing loosely on the pallavī from the composition 'vātāpi gaṇapatiṃ'), and each phrase/segment was represented as a directed graph using the python library networkx; the notes were represented as vertices, and vertices belonging to the same phrase were connected in order. For visualisation, vertices were plotted such that the y-coordinates represent their relative pitch positions. To get a connected graph, some vertices representing the same notes and sharing the local ascent–descent context in different phrases were merged. To enable splicing of phrases, the ending note of each phrase was connected to the beginning note of every phrase using virtual edges. This was subject to some constraints to avoid an excessively dense graph: only notes within a specified vertical distance from each other were connected with virtual edges. If a vertex simultaneously represented the beginning note of a phrase and the ending note of another phrase, it was labelled as a 'mixed note' (Fig. 1).

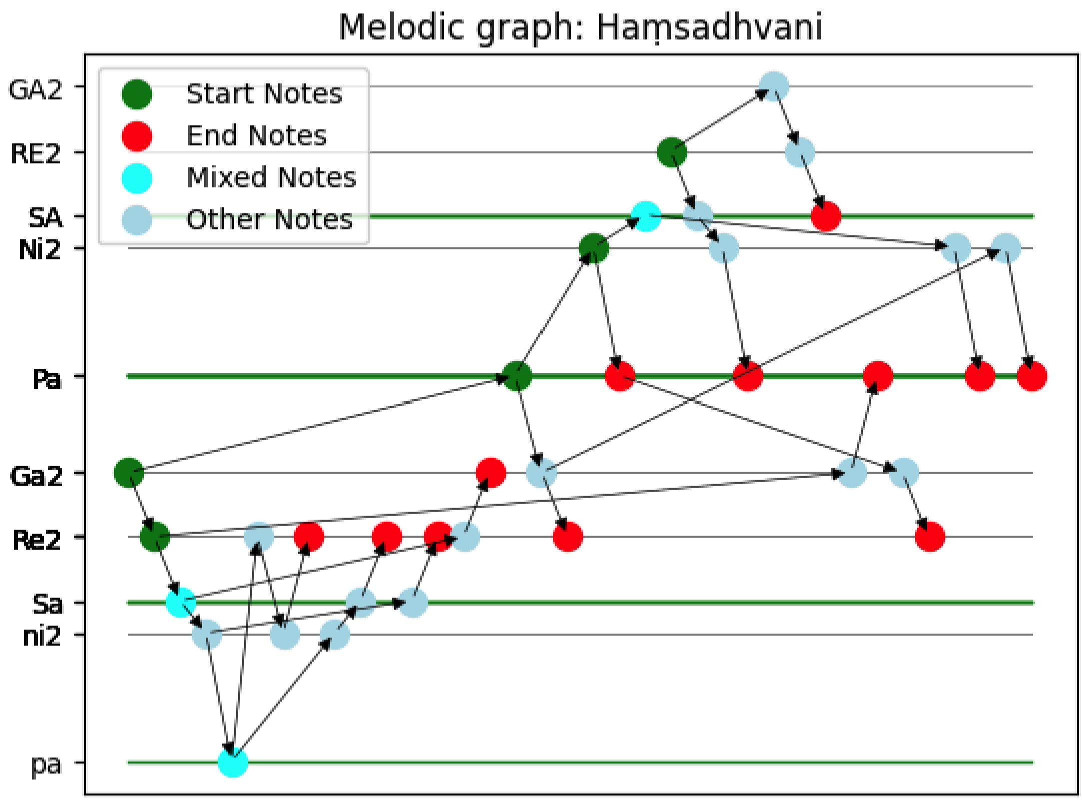

Fig. 1. Melodic graph generated using short note sequences.

New phrases were generated by tracing the path along the graph between pairs of notes (Fig. 2). Double notes occurring as a result of virtual edges connecting identical notes were removed. Duplicates/redundant phrases and very large phrases (those exceeding a threshold length) were also removed. From the remaining phrases, a random sample of 10 phrases was drawn. Another set of 10 sequences was generated by randomly sampling the notes available in the raga. The python code used in this study is available on request.

JOSHA
Journal of Science, Humanities and Arts
August 2025 Volume 12, Issue 4

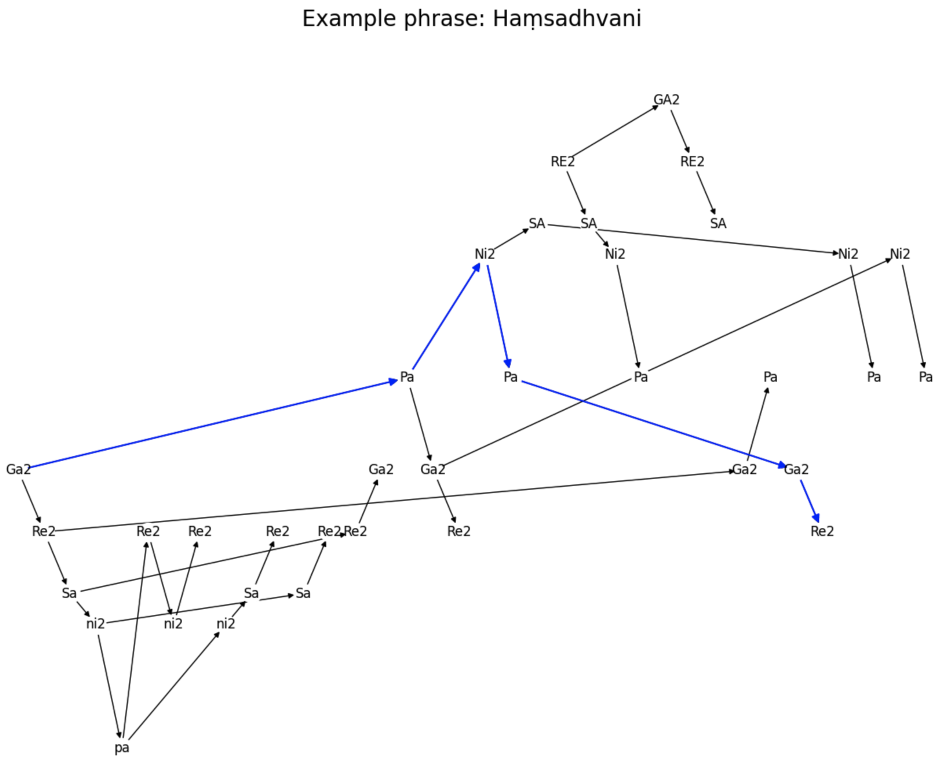

Fig. 2. A path (blue) traced along the graph gives a new phrase.

The generated phrases were then evaluated in comparison with random note sequences by 8 musicians trained in Carnatic music. The phrases—both generated phrases and random sequences—were presented in random order, and the musicians then evaluated them within the context of the raga and rated them on a scale of 1-10, with 1 representing no fit, and 10 representing a good fit of the phrase within the raga. The ratings for generated phrases and random phrases were averaged separately over the responses. A 2-tailed t-test was used to evaluate statistical significance for the difference in average scores between the 2 groups.

Results

The original phrases received an average rating of 7.6/10. The generated phrases received a significantly higher average rating than the random sequences (average rating of graph-generated sequences: 8.2/10; vs. random sequences: 3.6/10; $p << 0.01$; Fig. 3).

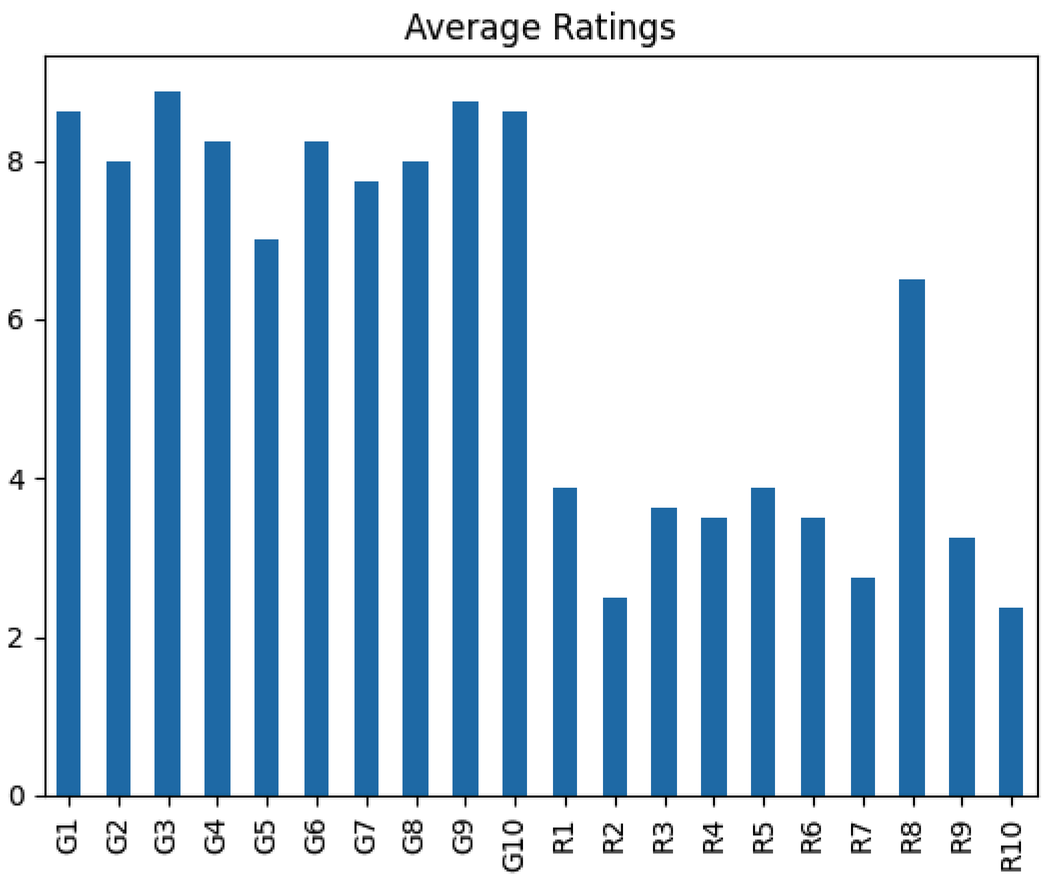

Fig. 3. Distribution of ratings for the generated phrases (G1-G10) and random sequences (R1-R10). The average rating for the G1-G10 group (mean = 8.2) is more than double the average rating for the R1-R10 group (mean = 3.6), and the difference is statistically significant ($p << 0.1$).

Discussion

Although melodic phrases are better modelled as pitch contours rather than as sequences of notes [2], Koduri et al. [6] have noted the potential for the use of symbolic scores for building more complex models of ragas. As noted above, finite-state models and Markov models (e.g., [4] and [5]) do not account for the history of transitions within a phrase. Thus, the current approach with directed graphs could be a promising improvement over these, allowing for highly context-specific modelling of ragas. Melodic graphs are intuitive and human-readable, and there is high explainability in its generative process. Because melodic improvisation is essentially based on the permutation and combination of notes [3]—along with ornamentation and emphasis through volume, duration, or timbre change—we can combine these phrases in new ways to generate new melodic material. The virtual edges in the graphs represent this splicing of phrases. In this regard, melodic graphs seem to have a large generative potential: the current graph created using just 16 very short segments is capable of generating hundreds of new, unique phrases. Jairazbhoy [1] noted that musicians develop 'musical habits' and evolve favourite phrases after repeatedly playing the same raga and tala. These habits can be thought of as familiar paths along the melodic graph with a high probability of treading upon. Melodic graphs may thus be able to capture stylistic variations among schools and even individual artists. These graphs may also help visualise individuals' personal, ever-evolving mental models of a raga. An interesting observation is that the entire pallavī of the Haṃsadhvani varṇam 'Jalajākṣa' can be described as paths on this graph, even though the graph is based on another composition in the same raga. This is particularly encouraging, as the graph seems to encapsulate some key elements of Haṃsadhvani. Melodic graphs may thus provide a robust theoretical model for ragas, describing composition and improvisation as graph-based probabilistic phenomena.

Limitations and Future Work

This study considered only a simple melodic graph built with just a few short sequences, which limits its generative power. The model can be improved with more sophisticated melodic graphs and associated methods. The ending note of each short phrase was connected (with virtual edges) to the beginning note of all short phrases within its defined neighbourhood. This effectively leads to indiscriminate connectivity without considering the local melodic context. More context-sensitive melodic graphs are needed to address this issue. In this study, notes were chosen arbitrarily as starting and ending notes of the generated phrases. However, in practice, the choice of these notes depends on the musical context [7]. More detailed work on phraseology can be taken up by modelling various ragas using melodic graphs. Improvisation is informed not just by the raga grammar but also by the aesthetic norms in practice [8]; it is also affected by the phraseology of other, similar ragas. Ornamentation or pitch, volume, and timbre dynamics form another indispensable part of many ragas' identity, which melodic graphs do not account for. More research is needed to accommodate these nuances for realistic computer-assisted composition, improvisation, and sound synthesis.

Conclusion

Melodic graphs are a novel, promising tool to represent ragas and model raga-based composition and improvisation. The melodic graph in the current study works well with Haṃsadhvani and has a large generative potential. Future work with other ragas and with more sophisticated melodic graphs is needed to fully explore their potential applications in music education, raga representation and modelling, and computer-assisted music generation.

References

1. Jairazbhoy, N.A.: The Rāgs of North Indian Music: Their Structure and Evolution. Faber and Faber, London (1971). https://eprints.soas.ac.uk/29748/1/10752720.pdf
2. Srinivasamurthy, A., Gulati, S., Caro Repetto, R., Serra, X.: Getting started on computational musicology and music information research: An Indian art music perspective. In: Rao, P., Murthy, H.S., Prasanna, S.R.M. (eds.) Indian art music: A computational perspective, pp. 3–38. Scheme for Promotion Academic and Research Collaboration, New Delhi (2023). http://hdl.handle.net/10230/56318
3. Rao, S., Rao, P.: An overview of Hindustani music in the context of computational musicology. Journal of New Music Research 43(1), 24–33 (2014). https://doi.org/10.1080/09298215.2013.831109
4. Das, D., Choudhury, M.: Finite state models for generation of Hindustani classical music. In: Proceedings of International Symposium on Frontiers of Research in Speech and Music, pp. 59–64 (2005). https://www.cs.cmu.edu/~dipanjan/pubs/frsm_gen.pdf
5. Varadharajan, J., Sridharan, G., Natarajan, V., Sridhar, R.: Automatic synthesis of notes based on Carnatic music raga characteristics. In: Kumar Kundu, M., Mohapatra, D., Konar, A., Chakraborty, A. (eds.) Advanced Computing, Networking and Informatics- Volume 1. Smart Innovation, Systems and Technologies, vol 27. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07353-8_17
6. Koduri, G. K., Miron, M., Serrà Julià, J., Serra, X.: Computational approaches for the understanding of melody in Carnatic music. In Klapuri, A., Leider, C. (eds.) ISMIR 2011: Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR 2011, pp 263–268. University of Miami, Miami (2011). https://repositori.upf.edu/handle/10230/22725
7. Krishna, T. M., Ishwar, V.: Karṇāṭik music: Svara, gamaka, phraseology and rāga identity. In: Serra, X., Rao, P., Murthy, H., Bozkurt, B. (eds.) Proceedings of the 2nd Comp-Music Workshop; Istanbul, Turkey, pp 12–18. Universitat Pompeu Fabra, Barcelona (2012). https://repositori.upf.edu/handle/10230/20494
8. Rao, S.: Culture specific music information processing: A perspective from Hindustani music. In: Serra, X., Rao, P., Murthy, H., Bozkurt, B. (eds.) Proceedings of the 2nd Comp-Music Workshop; Istanbul, Turkey, pp 5–11.

Universitat Pompeu Fabra, Barcelona (2012). https://repositori.upf.edu/handle/10230/20437

About the Authors

Bhushan Kumar Agarwal is a musicologist and educator exploring raga-based music—especially Dhrupad—through composition, teaching, and interdisciplinary research. With academic training in Biotechnology and Process Engineering, his current work combines digital signal processing, machine learning, and mathematical modelling with insights from performance and media studies. He studies gamakas, raga structures, and note sequences as phenomena that are embodied, shaped by history, and characterized by improvisation and interaction, aiming to develop nuanced computational representations. He's also interested in how traditional music shapes — and is shaped by — culture, society, and politics today.