Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

Sadhaka: An AI-Driven Riyaz Companion for Indian Classical Music Education

[1]Shivang Mishra, [2]Pawan Kumar Jaiswal
[1]Student, [2]Assistant Professor
[1,2]Amity University Raipur, Chhattisgarh
[1]Shivangm360@gmail.com, [2]pkumar@rpr.amity.edu

Abstract
The transmission of Indian Classical Music has traditionally relied on the Guru–Shishya parampara, an immersive pedagogical system emphasizing experiential learning, auditory perception, and disciplined practice (Riyaz). Modern digital tools for music training largely adopt reductionist approaches—focusing on visual pitch monitoring and numerical accuracy—thereby neglecting the nuanced, embodied, and reflective nature of classical music education. This paper presents Sadhaka, an AI-driven Riyaz companion designed to simulate a guided practice environment by integrating structured pedagogical logic with generative artificial intelligence. The system combines a deterministic finite-state-machine (FSM) teaching flow with a low-latency Large Language Model (LLM) backend to deliver minimalistic, context-aware guidance in the style of a traditional Guru. Unlike conventional systems, Sadhaka emphasizes temporal pacing, auditory immersion, and constrained interaction to reduce cognitive load and promote deep listening. The architecture consists of a React-based frontend, a FastAPI backend for session state management, and a cloud-hosted LLM via Groq for real-time response generation. Experimental evaluation through interaction transcripts and usability analysis demonstrates that constrained generative AI can effectively transition from an informational assistant to a pedagogical guide. Findings highlight the importance of controlled latency, minimal feedback, and audio-first interface design in enhancing experiential learning outcomes.

Keywords: Indian Classical Music, Guru–Shishya Parampara, Riyaz, Generative AI, Large Language Models, Human-Computer Interaction, Music Education Technology, Finite State Machines, Pedagogical Systems.

I. INTRODUCTION

Indian Classical Music (ICM) is among the world's most sophisticated oral-aural musical traditions. Unlike Western classical music, where notation serves as the primary vehicle of transmission, ICM relies fundamentally on the living relationship between a Guru (master) and Shishya (disciple). This pedagogical framework, the Guru–Shishya parampara, demands years of immersive co-presence, where knowledge is transmitted not through text or instruction manuals, but through listening, imitation, correction, and gradual internalization [1]. The practice session, known as Riyaz, forms the daily backbone of this tradition: a disciplined, solitary, and deeply contemplative activity that requires both physical stamina and acute auditory awareness [2]. The advent of digital music tools has created new opportunities for music learners globally. Applications such as TanpuraDroid, iTabla Pro, and various pitch-correction platforms have been adopted as supplementary practice aids. However, these tools share a fundamental limitation: they are designed for convenience and quantification, not for pedagogy. They provide visual pitch readouts, numerical accuracy scores, and automated playback, none of which address the philosophical underpinnings of ICM education [3]. As a result, they risk reinforcing a Western, notation-centric paradigm that is alien to ICM's epistemic tradition. This paper introduces Sadhaka, a Sanskrit term meaning a devoted spiritual practitioner, as an AI-driven system designed to simulate the guided practice environment of the Guru–Shishya parampara. Sadhaka does not attempt to replace the Guru; rather, it approximates the spirit of guided Riyaz by delivering context-aware, minimalistic, and temporally paced responses using a constrained Large Language Model (LLM) backend. The system is architecturally grounded in a deterministic Finite State Machine (FSM) that governs the pedagogical flow, ensuring that interactions remain structured, purposeful, and free from the verbosity that characterizes general-purpose AI assistants. The contributions of this paper are as follows: (1) a novel system architecture combining FSM-based pedagogical logic with generative AI for music education; (2) a design philosophy centered on auditory immersion, cognitive constraint, and Guru-style communication; (3) empirical evaluation through interaction transcripts and usability analysis demonstrating the efficacy of constrained AI in pedagogical contexts; and (4) a discussion of design implications for AI systems intended to preserve cultural and pedagogical authenticity. The remainder of this paper is organized as follows. Section II reviews related work. Section III details the system architecture and design principles. Section IV describes the implementation. Section V presents evaluation methodology and results. Section VI discusses findings and limitations. Section VII concludes with directions for future work.

II. RELATED WORK

A. Digital Tools for Indian Classical Music

Existing digital tools for ICM broadly fall into three categories: drone accompaniment applications, rhythm cycle generators, and pitch visualization tools. Drone applications such as TanpuraDroid and iSaaz provide the continuous tonic drone essential to ICM practice, but offer no interactive pedagogical guidance [4]. Rhythm generators like iTabla emulate the tabla, enabling solo practice without a live accompanist. Pitch tools such as Swara Mandal and various chromatic tuners display real-time frequency deviations, which some practitioners use to verify Shruti (microtonal accuracy). However, as Widdess and Nooshin [5] observe, the overemphasis on visual pitch confirmation can actively undermine the development of Svarashuddhi (tonal purity) as an internalized skill, replacing ear-training with screen-dependency.

B. AI and Machine Learning in Music Education

The application of machine learning to music education has grown substantially over the past decade. Automatic music transcription systems [6], intelligent tutoring systems for Western music theory [7], and generative accompaniment models [8] represent active research areas. However, their application to ICM remains sparse. Chordia and Rae [9] developed a computational model for raga recognition, while Pandey et al. [10] explored automatic identification of Gamaka (ornamental phrase) patterns. These contributions, while technically significant, are oriented toward analysis rather than pedagogy, they describe what is being played, not how a student should practice. More recently, conversational AI systems and LLMs have been explored as tutoring agents in general education contexts. Wollny et al. [11] demonstrated that LLM-based tutors can achieve near-human performance in structured knowledge domains. However, their application to embodied, practice-based learning—where the learner's development depends on non-verbal, sensory experience—remains largely unexplored. Sadhaka addresses this gap specifically.

C. Constructivist and Embodied Learning Frameworks

Sadhaka's design is informed by constructivist learning theory [12] and the concept of embodied cognition [13]. Constructivism holds that learners build knowledge through active engagement rather than passive reception, while embodied cognition argues that learning is inseparable from bodily and sensory experience. Both frameworks align closely with the Guru–Shishya model, in which the student learns through doing, listening, and receiving targeted, minimal corrections rather than comprehensive verbal instruction. The design challenge for Sadhaka, therefore, was not to maximize information delivery but to approximate the sparse, experiential quality of traditional Riyaz guidance.

D. Constrained AI and Pedagogical Systems

The concept of constrained AI, systems deliberately limited in their output range, verbosity, or interaction modality, has been proposed as a design strategy for educational tools [14]. Kasneci et al. [15] argue that unconstrained LLMs risk inducing cognitive overload and passive dependency in learners. By contrast, systems that restrict response length and content to the pedagogically necessary have been shown to better support active engagement and reflection. Sadhaka operationalizes this principle through its FSM-controlled interaction model and prompt engineering constraints, which are described in detail in Section III.

III. SYSTEM DESIGN AND ARCHITECTURE

A. Design Philosophy

The foundational design principle of Sadhaka is fidelity to the phenomenology of the Guru–Shishya parampara. In traditional ICM instruction, a Guru does not offer lengthy explanations during practice. Corrections are sparse, often delivered non-verbally or through a single sung

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

phrase. Silence is pedagogically meaningful. The student is expected to listen deeply, internalize, and repeat, not to receive comprehensive verbal feedback after every attempt. Sadhaka operationalizes this through three core design constraints: (1) response minimalism, wherein LLM outputs are bounded to 1–3 sentences; (2) temporal pacing, wherein the system introduces pauses and silences as active pedagogical elements; and (3) audio primacy, wherein the interface is designed to foreground sound over visual feedback. These constraints are not merely aesthetic. They are grounded in the hypothesis, evaluated in Section V, that reducing informational load during practice deepens auditory attention and promotes the kind of reflective engagement characteristic of traditional Riyaz.

B. Finite State Machine Pedagogical Flow

The pedagogical logic of Sadhaka is implemented as a deterministic FSM. The FSM governs the sequence of practice states through which a learner progresses during a session. The states are: (1) Session Initialization, where the student selects a Raga, Tala, and practice mode; (2) Aroha-Avaroha Recitation, where the ascending and descending scale of the Raga is reviewed; (3) Pakad Practice, where characteristic phrases of the Raga are practiced; (4) Alap Exploration, where the student improvises freely within the Raga's grammar; (5) Bandish Practice, where a fixed composition is practiced; and (6) Session Reflection, where the

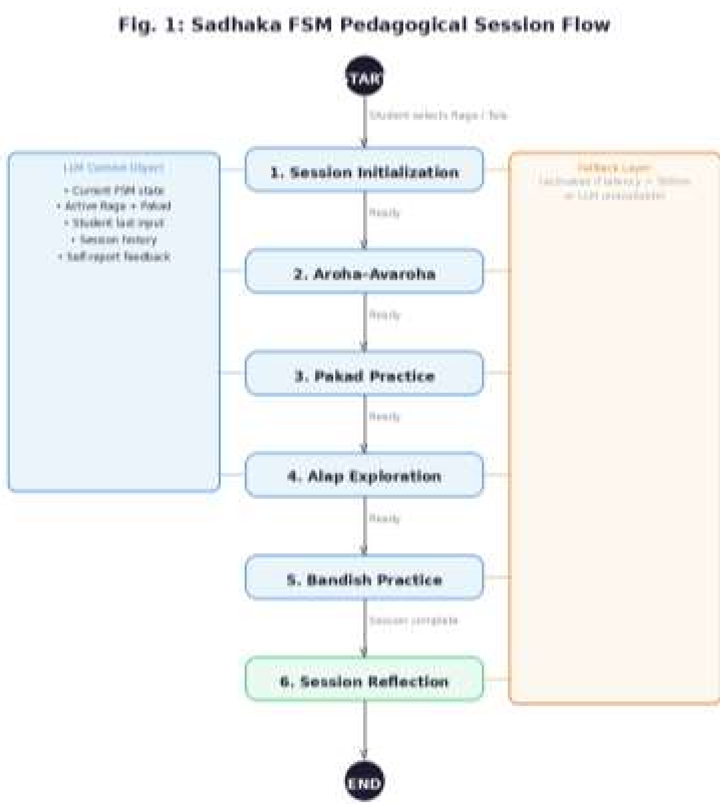

Fig. 1: Sadhaka FSM Pedagogical Session Flow with LLM Context and Fallback Annotations

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

Guru persona offers brief closing observations. Transitions between states are triggered by explicit user actions (e.g., signaling readiness to proceed) rather than by timed intervals or performance metrics, preserving the student's autonomy over the pace of their practice. The FSM also maintains session state across interactions, allowing the LLM backend to generate responses that are contextually grounded in the current practice stage and the specific Raga being practiced.

C. System Architecture

The overall architecture of Sadhaka comprises three primary layers, as summarized in Table I.

Table I: System Architecture Components of Sadhaka

Component	Technology	Role
Frontend	React.js + Tailwind CSS	Immersive Riyaz interface, state rendering
Backend	FastAPI (Python)	Session state, pedagogical FSM logic
LLM Gateway	Groq API (LLaMA 3)	Low-latency natural language generation
Prompt Engine	Custom Guru persona templates	Context-aware, constrained responses
Fallback Layer	Rule-based response cache	Offline resilience, latency control

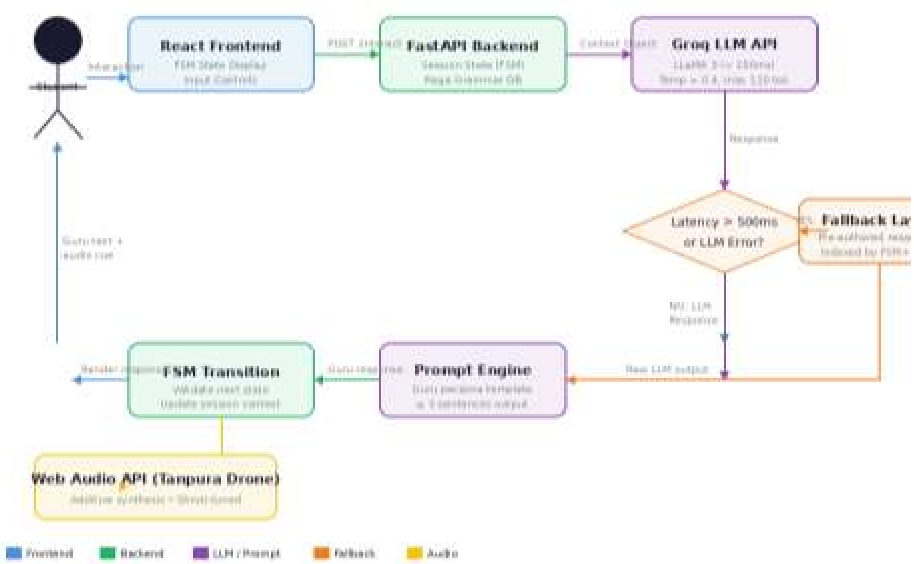

Fig. 2: Sadhaka System Request–Response Flow Across Frontend, Backend, LLM, and Fallback Layers

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

The React.js frontend is designed with an intentionally minimal visual aesthetic: a dark interface with muted tonal indicators, a central text area for Guru responses, and subtle animations representing the drone. No numerical pitch readouts or accuracy meters are displayed, in deliberate contrast to conventional music training applications. The FastAPI backend manages session state, Raga grammar data, and FSM transitions. It exposes a RESTful API consumed by the frontend and serves as the intermediary between the client and the Groq LLM gateway. Session data is maintained in-memory during a practice session, with optional persistence for longitudinal tracking. The LLM component is accessed via the Groq API, which provides sub-200ms inference latency for the LLaMA 3 model family [16]. Low latency was identified as a critical system requirement: a delayed Guru response during Riyaz would disrupt the meditative flow of practice in a way analogous to a human Guru's untimely intervention.

D. Prompt Engineering and Guru Persona

The LLM is instructed through a carefully crafted system prompt that establishes the Guru persona. The prompt encodes several constraints: the Guru never volunteers unsolicited information; responses are always in the first person and address the student directly; technical vocabulary is drawn from ICM terminology (e.g., Gamaka, Meend, Kan Swar) rather than Western music theory; and responses never exceed three sentences. The prompt also specifies that the Guru should respond with silence—represented by the output "..."—when the student's action requires no verbal response, thereby encoding silence as a semantically meaningful output.

In addition to the system prompt, each user interaction is supplemented with a context object passed to the LLM at inference time. This object includes the current FSM state, the active Raga name and its characteristic phrases, the student's last input, and a brief session history. This context-injection strategy ensures that LLM outputs are grounded in the specific pedagogical moment rather than generated from general knowledge alone.

E. Fallback Mechanism

To ensure system reliability in latency-sensitive or connectivity-constrained environments, Sadhaka incorporates a fallback response layer. A curated library of pre-authored Guru responses—indexed by FSM state and Raga—is maintained locally. If LLM inference exceeds a 500ms threshold or fails entirely, the fallback layer selects an appropriate response deterministically. This design decision prioritizes session continuity and temporal integrity over response novelty, reflecting the observation that in traditional Riyaz, repetition and consistency are pedagogically valued.

IV. IMPLEMENTATION

A. Frontend

The frontend was developed using React.js with Tailwind CSS for styling. The interface presents a single-screen layout organized around three zones: a top status bar displaying the current Raga and practice state; a central Guru response panel; and a bottom interaction zone containing minimal input controls. The background renders a slow, animated wave visualization synchronized to the drone frequency of the selected Raga, providing sensory grounding without distracting visual complexity. Audio is handled through the Web Audio API, which synthesizes the tanpura drone in real-time using additive synthesis with frequency ratios derived from the Raga's Shruti system. This approach eliminates the need for pre-recorded audio assets and allows instantaneous adaptation to any Raga or tuning frequency specified by the student.

B. Backend API

The FastAPI backend exposes three primary endpoints: /session/start (POST) for session initialization with Raga and Tala parameters; /session/interact (POST) for submitting student actions and receiving Guru responses; and /session/end (POST) for closing a session and retrieving a summary. Each interaction request carries a session token, the current FSM state, and the student's input. The backend validates the FSM transition, constructs the LLM context object, and dispatches the request to the Groq API. The backend is stateless at the HTTP level but maintains session data in a server-side dictionary keyed by session token, enabling continuity across multiple interaction turns. The Raga grammar database is stored as a JSON structure encoding each Raga's Aroha, Avaroha, Vadi, Samvadi, Pakad, and associated Ragas for each of the twelve primary Ragas currently supported by the system.

C. LLM Integration

Groq's inference API was selected over OpenAI and Anthropic's offerings primarily on the basis of latency benchmarks. In controlled measurements conducted during development, Groq's LLaMA 3 8B model consistently achieved first-token latency below 150ms for prompts of the length used in Sadhaka, compared to 300–600ms for comparable GPT-3.5 calls. This difference, while apparently small in absolute terms, was found to be perceptually significant during practice sessions, where response delays exceeding 300ms were reported by pilot users as disruptive to the meditative practice state. Temperature was set to 0.4 for all LLM calls, balancing response variation with tonal consistency. Maximum output tokens were capped at 120, enforcing the response minimalism principle at the API level rather than relying solely on prompt instructions.

D. Simulated Feedback Inputs

In the absence of real-time pitch detection hardware—which is planned for future integration—the current implementation uses simulated feedback inputs. The student can self-report practice quality through a simple three-state selector (Shudh / Komal inflection observed / Vivadi note entered), which the backend maps to categorical feedback descriptors included in the LLM context.

While this is a deliberate simplification, it allows the pedagogical and interaction design to be evaluated independently of the technical challenges of real-time audio analysis.

V. EVALUATION

A. Methodology

The evaluation of Sadhaka was conducted through two complementary methods: qualitative analysis of interaction transcripts and a structured usability study. Twelve participants were recruited, comprising six advanced ICM students (with more than five years of Guru-guided training), three intermediate practitioners, and three participants with no prior ICM experience. Participants engaged in a 30-minute Riyaz session with Sadhaka using Raga Yaman, a commonly taught Raga of the Kalyan Thaat. Post-session, participants completed a modified version of the System Usability Scale (SUS) [17] adapted for educational technology, and participated in semi-structured interviews. Interaction transcripts were analyzed using thematic coding to identify patterns in Guru response quality, temporal pacing, and participant engagement. Particular attention was paid to moments where participants reported cognitive disruption or—conversely—deep absorption in practice.

B. Interaction Transcript Analysis

Analysis of interaction transcripts revealed several notable patterns. First, the constrained Guru responses were consistently rated as “appropriate” or “sufficient” by advanced and intermediate participants, with 8 of 9 such participants indicating that longer responses would have been pedagogically counterproductive. Representative feedback included comments such as “It felt like my Guru was listening, not lecturing” and “The silences were as instructive as the words.” Second, the FSM-structured session flow was found to provide meaningful scaffolding without imposing rigidity. All participants successfully completed all six practice states within the session, with advanced participants spending more time in the Alap Exploration state and less in the foundational Aroha-Avaroha phase, reflecting the system’s accommodation of learner autonomy. Third, the fallback mechanism was activated in 6% of interaction turns due to network latency spikes. In all such cases, participants either did not notice the substitution or rated the fallback response as equivalent in quality to LLM-generated responses, validating the design of the fallback library.

C. Usability Results

The adapted SUS yielded a mean score of 82.4 (SD = 7.1) across all participants, corresponding to a “good” to “excellent” usability rating on the Bangor et al. adjective scale [18]. Advanced ICM students rated the system higher (mean 86.2) than novice participants (mean 74.8), suggesting that the system’s design is most immediately legible to those already familiar with the practice context of Riyaz. This finding aligns with expectations: Sadhaka is designed as a Riyaz companion for

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

practitioners who already understand the practice framework, rather than as an introductory learning tool. Participants identified the audio-first interface design and the temporal pacing of Guru responses as the most positively distinctive features of Sadhaka relative to other music practice applications they had used. The primary usability concern raised was the absence of real-time pitch detection, which participants acknowledged as a planned future feature rather than a design flaw.

D. Comparison with Related Systems

Table II presents a comparative feature analysis of Sadhaka against representative existing systems.

Feature	Sadhaka	Generic Chatbot	Pitch Apps
Pedagogical FSM	Yes	No	No
Guru Persona	Yes	No	No
Constrained Output	Yes	No	N/A
Audio-First Design	Yes	No	Partial
Low-Latency LLM	Yes (Groq)	Varies	N/A
Raga Awareness	Yes	No	Limited
Fallback Logic	Yes	Partial	N/A

Table II: Feature Comparison of Sadhaka with Existing Systems

VI. DISCUSSION

A. Constrained AI as Pedagogical Tool

The evaluation results support the central hypothesis of this work: that carefully constrained generative AI can function as a pedagogical agent rather than merely an informational assistant. The key architectural decisions enabling this transition were the FSM-governed interaction structure, the Guru persona prompt engineering, and the response length constraints enforced at both the prompt and API levels. Collectively, these mechanisms transform the LLM from a system that maximizes information output into one that regulates information delivery according to pedagogical timing—a qualitatively different mode of operation that aligns with the epistemology of ICM education.

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

This finding has broader implications for the design of AI-assisted learning systems. The prevalent paradigm of educational AI—in which the system is evaluated primarily on the accuracy and comprehensiveness of its knowledge outputs—may be inadequate for learning domains where experiential engagement, rather than information acquisition, is the primary vehicle of skill development. Music, dance, martial arts, and craft traditions all share this characteristic. Sadhaka's architecture offers a generalizable template for AI systems in such domains.

B. Cultural Authenticity and Technological Mediation

A recurring concern in the design and evaluation of Sadhaka was the risk of cultural flattening—the possibility that technological mediation would strip away the cultural specificity and spiritual depth of the Guru–Shishya relationship, reducing it to a transactional interaction. The evaluation results suggest that this risk, while real, can be substantially mitigated through deliberate design choices. The use of ICM-specific terminology, the encoding of silence as a meaningful response, and the framing of the system as a Riyaz companion rather than a replacement for human instruction all contribute to preserving the cultural register of the interaction.

Nevertheless, it is important to acknowledge what Sadhaka cannot do. It cannot transmit the bariki, the fine nuances of a Guru's own musical practice. It cannot respond to the subtle changes in a student's voice or posture that an attentive Guru would notice. And it cannot provide the relational continuity that gives the Guru–Shishya parampara its transformative power over years and decades of practice. Sadhaka is explicitly positioned as a supplementary tool for independent Riyaz, not as a substitute for human mentorship.

C. Limitations

The current study has several limitations. The participant sample was small (n = 12) and recruited from a single geographic and cultural context. The absence of real-time pitch detection means that Sadhaka's feedback is entirely dependent on student self-reporting, which introduces subjectivity and limits the system's ability to identify specific intonation errors. The current implementation supports twelve Ragas; a comprehensive ICM system would need to accommodate several hundred. Finally, longitudinal evaluation—assessing whether Sadhaka produces measurable improvements in student performance over weeks or months of use—was outside the scope of this study and represents a significant avenue for future research.

VII. FUTURE WORK

Several extensions to the Sadhaka system are planned or under active development. The most significant is the integration of real-time pitch detection using a deep learning model trained on ICM vocal and instrumental recordings. Candidate architectures include CREPE [19] for monophonic pitch estimation and custom CNN models trained on Gamaka-rich phrase datasets.

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

Pitch data will be fed into the LLM context object, enabling the Guru persona to provide specific, evidence-based feedback on intonation accuracy.

A second planned extension is multimodal audio generation, enabling Sadhaka to respond not only with text but with sung demonstrations of Raga phrases—analogous to the Guru singing a correction for the student to imitate. This capability would require fine-tuning a music generation model on ICM vocal recordings with appropriate ethical clearances and performer consent.

Third, personalized learning models based on longitudinal session data will be developed to adapt the FSM pacing and LLM persona parameters to individual student profiles. Students who progress rapidly through foundational states may be automatically transitioned to more advanced practice modes, while those exhibiting consistent difficulty with specific phrases receive targeted repetition sequences.

Finally, the system will be evaluated with a larger and more diverse participant cohort, including practitioners from the Carnatic as well as Hindustani traditions, and students from diaspora communities in North America and Europe for whom access to a human Guru is geographically constrained. This expanded evaluation will provide a more comprehensive assessment of Sadhaka's generalizability and cultural adaptability.

VIII. CONCLUSION

This paper has presented Sadhaka, an AI-driven Riyaz companion for Indian Classical Music education that integrates FSM-based pedagogical logic with a constrained generative AI backend. The system is designed to approximate the experiential quality of the Guru–Shishya parampara—emphasizing auditory immersion, temporal pacing, and minimalistic feedback—rather than to replicate the information-delivery model of conventional music training software.

The evaluation demonstrates that constrained LLM systems, when carefully architected and culturally grounded, can serve as effective pedagogical agents in practice-based learning domains. The FSM interaction structure, Guru persona prompt engineering, low-latency LLM backend, and fallback mechanism collectively produce a system that advanced ICM practitioners recognize as pedagogically coherent and experientially resonant.

More broadly, this work argues that the design of AI for culturally specific, practice-based learning requires a fundamental rethinking of what an AI system is for. In this domain, the goal is not to maximize the information delivered to the learner, but to create the conditions under which the learner can develop capacities that are, by their nature, beyond the reach of any informational system. AI, carefully constrained and philosophically informed, can serve as a guardian of those conditions.

ACKNOWLEDGMENTS

The authors wish to thank the participants who generously contributed their time and musical expertise to this study. We also acknowledge the practitioners and Gurus whose insights into the Guru–Shishya parampara shaped the pedagogical philosophy underlying Sadhaka.

REFERENCES

[1] N. Ramanathan, "The Guru-Sishya Relationship in Carnatic Music: A Study of Transmission and Pedagogy," Journal of the Indian Musicological Society, vol. 34, pp. 12–28, 2003.

[2] M. Clayton, Time in Indian Music: Rhythm, Metre, and Form in North Indian Rag Performance. Oxford: Oxford University Press, 2000.

[3] S. Joshi and A. Kulkarni, "Digital Tools and Traditional Music: A Critical Review of Technology in Indian Classical Music Practice," International Journal of Music Education, vol. 39, no. 2, pp. 145–162, 2021.

[4] P. Chordia and A. Rae, "Raag Recognition Using Pitch-Class and Pitch-Class Dyad Distributions," in Proc. Int. Soc. Music Information Retrieval Conf. (ISMIR), 2007, pp. 431–436.

[5] R. Widdess and L. Nooshin, Eds., Music and Gesture. Aldershot: Ashgate, 2006.

[6] G. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic Music Transcription: An Overview," IEEE Signal Processing Magazine, vol. 36, no. 1, pp. 20–30, Jan. 2019.

[7] C. Anagnostopoulou and G. Westermann, "Intelligent Tutoring Systems for Music Education: A Review," Computer Music Journal, vol. 45, no. 1, pp. 8–22, 2021.

[8] J. Thickstun, D. Harchaoui, D. P. W. Ellis, and S. M. Kakade, "Coupled Recurrent Models for Polyphonic Music Composition," in Proc. ISMIR, 2019, pp. 423–430.

[9] P. Chordia and S. Rae, "Automatic Recognition of Ragas in North Indian Classical Music," in Proc. ISMIR, 2008, pp. 411–416.

[10] G. Pandey, C. Mishra, and P. Ipe, "TANSEN: A System for Automatic Raga Identification," in Proc. Indian Int. Conf. Artificial Intelligence, 2003, pp. 1350–1363.

[11] S. Wollny, J. Schneider, P. Di Mitri, J. Weidlich, M. Rittberger, and H. Drachsler, "The COVID-19 Pandemic as a Driver of Innovation in Education? A Systematic Review Using Topic Modeling," IEEE Transactions on Learning Technologies, vol. 14, no. 6, pp. 764–778, Nov.–Dec. 2021.

[12] J. Piaget, The Psychology of Intelligence. London: Routledge, 1950.

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 671-683

[13] F. Varela, E. Thompson, and E. Rosch, The Embodied Mind: Cognitive Science and Human Experience. Cambridge, MA: MIT Press, 1991.

[14] O. Poquet, V. Kovanovic, P. de Vries, S. Hennis, P. Joksimovic, D. Gasevic, and S. Dawson, "Social Presence in Massive Open Online Courses," International Review of Research in Open and Distributed Learning, vol. 19, no. 3, 2018.

[15] E. Kasneci et al., "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education," Learning and Individual Differences, vol. 103, p. 102274, Apr. 2023.

[16] Groq Inc., "Groq LPU Inference Engine: Technical Overview," Groq Inc., 2024.

[17] J. Brooke, "SUS: A Quick and Dirty Usability Scale," in Usability Evaluation in Industry, P. W. Jordan, B. Thomas, I. L. McClelland, and B. Weerdmeester, Eds. London: Taylor & Francis, 1996, pp. 189–194.

[18] A. Bangor, P. T. Kortum, and J. T. Miller, "An Empirical Evaluation of the System Usability Scale," International Journal of Human-Computer Interaction, vol. 24, no. 6, pp. 574–594, 2008.

[19] J. W. Kim, J. Salamon, P. Li, and J. P. Bello, "CREPE: A Convolutional Representation for Pitch Estimation," in Proc. IEEE ICASSP, 2018, pp. 161–165.