



Review Article

FROM MANUSCRIPT TO MODEL: LARGE LANGUAGE MODELS AS TOOLS FOR AYURVEDIC
KNOWLEDGE PRESERVATION AND ACCESS

Nikita Sarkar^{1*}, Neha Kumawat²

¹Assistant Professor, Dept. of Samhita Siddhanta, Raghunath Ayurved Mahavidyalaya & Hospital, Contai, East Midnapore, West Bengal.

²PG Scholar, Dept. of Shalya Tantra, Shri Krishna Government Ayurvedic College and Hospital, Kurukshetra, Haryana, India.

Article info

Article History:

Received: 26-06-2026

Accepted: 22-07-2026

Published: 15-08-2026

KEYWORDS:

Large Language Models (LLMs), Artificial Intelligence, Domain-Specialized Language Models, Retrieval-Augmented Generation (RAG), Sanskrit Knowledge Digitization.

ABSTRACT

Ayurveda, one of the world's oldest continuously practiced medical systems, encodes centuries of clinical knowledge in Sanskrit and regional-language texts that remain largely inaccessible to modern computational tools. The emergence of large language models (LLMs) has opened new possibilities for digitizing, interpreting, and operationalizing this knowledge base at scale. This review synthesizes current research on LLMs applied to Ayurveda, covering domain-specialized models such as AyurParam and AyurGPT, retrieval-augmented generation, machine-learning approaches to *Prakriti* classification, and digitization of classical texts. It examines key challenges, including hallucination risk, the linguistic complexity of Sanskrit, data standardization gaps, and unresolved ethical and regulatory questions, and contributes a small original case study comparing a general-purpose LLM's behavior against documented domain-specialized systems. The review concludes that domain-adapted LLMs show measurable gains over general-purpose models on Ayurveda-specific tasks, but large-scale clinical validation and governance frameworks remain necessary before such systems can be safely integrated into practice.

INTRODUCTION

Ayurveda is a traditional system of medicine that originated in the Indian subcontinent more than two thousand years ago and remains in active use today, with a large share of India's population relying on it either exclusively or alongside conventional biomedicine. The system rests on foundational concepts such as the *Tridosha* theory (*Vata*, *Pitta*, and *Kapha*), *Prakriti* (individual constitution), and *Samprapti* (disease pathogenesis), described across a small set of classical Sanskrit compendia that together define what any AI system claiming Ayurvedic grounding must ultimately be validated against. [1,2,3,4]

Large language models have progressed rapidly since the introduction of instruction-tuning techniques that teach models to follow natural-

language directions reliably, moving from narrow task-specific systems toward broadly capable conversational assistants. However, general-purpose LLMs trained predominantly on Western biomedical and English-language corpora tend to perform poorly on the specialized vocabulary, philosophical framing, and linguistic structure of Ayurvedic texts, a gap that mirrors a wider pattern in which domain specialization is increasingly seen as necessary for reliable, explainable performance on expert tasks. [5,6,7]

This gap has motivated a growing body of research adapting LLMs specifically for Ayurveda, spanning domain fine-tuning, retrieval-augmented generation, and hybrid machine-learning pipelines for diagnostic tasks such as *Prakriti* classification. The motivation is not purely academic: India's Ministry of AYUSH has been building digital health infrastructure for several years, and more recent national programs such as Aikosh have explicitly framed large-scale AI adoption as a mechanism for validating and modernizing traditional medicine at population scale. [8]

Access this article online	
Quick Response Code	
	https://doi.org/10.47070/ijapr.v14i8.4340
Published by Mahadev Publications (Regd.) publication licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)	

This convergence of policy interest and rapid LLM progress make the present moment a useful point at which to take stock of what has been achieved and what remains unresolved. This paper reviews the current state of research at this intersection, surveys the principal application areas that have emerged in the literature, contributes a small original case study, and discusses the technical, linguistic, and ethical challenges that must be addressed before such systems are responsibly deployed in clinical or educational settings.

AIMS AND OBJECTIVES

The aim of this review is to evaluate the application of large language models (LLMs) in Ayurveda and assess their potential for preserving and improving access to Ayurvedic knowledge. Specifically, this review aims to:

1. Summarize major applications of LLMs in Ayurveda, including domain-specific models, retrieval-augmented generation, *Prakriti* classification, digitization of classical texts, and clinical decision support.
2. Examine key technical, linguistic, and ethical challenges.
3. Compare the performance of general-purpose and domain-specific LLMs through a case study.
4. Discuss their responsible integration into Ayurvedic education, research, and clinical practice.

METHODOLOGY

This paper follows a narrative review approach rather than a formal systematic review. Sources were identified through targeted searches of academic databases and preprint repositories, including arXiv, IEEE Xplore, and journals covering Ayurveda and integrative medicine, supplemented by recent coverage of national digital-health initiatives relevant to Ayurveda. Priority was given to work published within the past two to three years, reflecting the pace of change in both LLM research generally and Ayurveda-specific AI research in particular, and the resulting synthesis is organized thematically around application areas and cross-cutting challenges rather than chronologically.

Background

Large Language Models

Large language models are transformer-based neural networks trained on massive text corpora to predict and generate coherent language. General-purpose variants demonstrate broad competence across many tasks but underperform in specialized domains requiring deep cultural, linguistic, or subject-matter expertise, a limitation that has motivated a broader research direction toward domain

specialization as a way to make large models more efficient, reliable, and explainable for narrow expert tasks rather than relying solely on raw scale. [7]

Ayurveda as a Knowledge Domain

Ayurvedic diagnosis and treatment planning rely on constructs with no direct biomedical equivalent, such as *Prakriti* classification and *Dosha*-imbalance identification, which trace back to the Charaka Samhita, Sushruta Samhita, Ashtanga Hridaya, and Kashyapa Samhita, the foundational texts that together form the classical basis for the field. These assessments traditionally depend heavily on individual practitioner expertise and are therefore prone to subjectivity and inter-practitioner variability, which has created interest in computational methods that could make such assessments more standardized and reproducible while remaining faithful to the classical sources from which Ayurvedic knowledge derives. [1,2,3,4,9]

Applications of LLMs in Ayurveda

Domain-Specialized Language Models

AyurParam-2.9B, introduced in late 2025, is a bilingual (English/Hindi) model fine-tuned from the PARAM-1 base model using classical texts, clinical guidance documents, and expert-annotated question-answer pairs, representing one of the first large-scale attempts to build a language model purpose-built for this domain rather than adapted from a general medical model. [10,11]

On BhashaBench-Ayur, a benchmark purpose-built for Ayurvedic knowledge, AyurParam outperformed larger general-purpose open-source models several times its size, illustrating a broader pattern in specialized-domain LLM research: targeted fine-tuning on a smaller, high-quality dataset can outperform raw parameter scale once a domain is sufficiently distinct from general web text. However, its own reported results reveal a meaningful gap between its two supported languages: accuracy on BhashaBench-Ayur reached 41.12 percent for English queries but only 38.04 percent for Hindi queries, a practically significant limitation given that Hindi is one of the primary languages in which Ayurvedic practitioners and patients actually communicate, and one the developers themselves attribute to insufficient Hindi-language representation in the training corpus. [10,12]

AyurParam is not the only domain-specific effort in this space. AyurGPT has similarly been presented as a fine-tuned medical language model aimed specifically at multilingual Ayurveda consultations, indicating that domain-specialized fine-tuning for Ayurveda is becoming an active area with more than one competing approach rather than a single isolated project. [9,13]

Retrieval-Augmented Generation

A parallel line of work addresses the reliability problem through retrieval-augmented generation (RAG) rather than, or in addition to, fine-tuning. AyuRAG combines a lightweight LLM with a retrieval layer that pulls relevant passages from a structured Ayurvedic knowledge base before generating a response, grounding its answers in traceable source material rather than relying solely on parametric memory. This approach is particularly valuable given how fragmented and inconsistently digitized Ayurvedic literature remains, since it allows a system to stay useful even before a fully comprehensive fine-tuning corpus exists.^[14]

Prakriti Classification and Diagnostic Support

A substantial body of machine-learning research, adjacent to but distinct from LLM research proper, has focused on automating *Prakriti* classification from questionnaire data and biometric signals. Early supervised approaches applying classifiers to constituent-balance prediction achieved reported accuracies in the 90-95 percent range, and subsequent work has extended this into more explicitly AI-branded frameworks describing *Prakriti* analysis as a convergence point between Ayurveda and modern technology.^[15,16]

More recent Ayurgenomics research has gone further, seeking a biological basis for constitution itself by reporting associations between specific gene variants and an individual's *Prakriti*-based disease predisposition, suggesting that *Prakriti* classification may eventually be groundable in molecular as well as behavioral data rather than questionnaire responses alone.^[17]

Related efforts have modernized *Nadi Pariksha* (pulse-based diagnosis) using Internet-of-Things sensor platforms for real-time patient monitoring and automated disease prediction, with comparable IoT-based approaches extended to broader Ayurvedic healthcare monitoring beyond pulse diagnosis alone. At the same time, a scoping review of *Prakriti* evaluation tools found considerable inconsistency in the psychometric properties of assessment instruments used across studies, underscoring that standardization remains an open problem even for the underlying human assessments that many of these AI systems are trained to replicate.^[18,19,20]

Digitization of Classical Texts

Because Ayurvedic knowledge is dispersed across Sanskrit manuscripts, regional-language commentaries, and modern clinical guidelines, natural language processing plays a foundational role in making this material machine-readable. Much of the currently digitized corpus available for this purpose

resides in general-purpose public archives, including the Internet Archive, the eGangotri digital library, and the National Digital Library of India, rather than in repositories built specifically for Ayurvedic or Indic-language content. National initiatives such as India's AIKosh program aim to provide standardized, large-scale digital infrastructure connecting practitioners, research institutions, and AI developers, addressing the data-fragmentation problem that has historically limited progress in this area.^[8,21,22,23]

Clinical Decision Support

Concrete, condition-specific applications are beginning to emerge beyond general proof-of-concept work. One machine-learning-based clinical decision support system was built specifically to stratify gestational diabetes mellitus and guide its management through Ayurvedic principles, combining patient data with structured Ayurvedic reasoning to produce individualized recommendations, illustrating how such tools are moving from broad diagnostic support toward narrower, clinically validated use cases.^[24]

Drug Discovery and Formulation Research

Machine-learning and network-analysis techniques have also been applied to the study of Ayurvedic herbal formulations. One study used supervised network analysis and machine learning to investigate potential anti-bacterial compounds embedded within classical Ayurvedic formulae, representing a concrete bridge between traditional pharmacology and modern computational drug discovery. Broader reviews of AI in Ayurvedic pharmaceuticals catalogue both such opportunities, including accelerated candidate identification, and persistent challenges, including limited standardization of herbal source material and the difficulty of validating traditional formulations against modern pharmacological evidence standards.^[25,26]

Evaluation and Benchmarking

A necessary companion to any application-focused research is a means of measuring progress. BhashaBench-Ayur represents an early but important step toward standardized assessment of Ayurveda-specific language understanding, covering objective question-answering and contextual reasoning about dosha imbalances in both English and Hindi. Researchers have separately argued that medical benchmarks for large language models more broadly need to be rethought to better reflect real clinical reasoning rather than narrow factual recall, a critique that applies directly to Ayurveda-specific benchmarking as well.^[12,27]

Systematic scoping reviews of AI use across complementary and alternative medicine, a category that includes Ayurveda, similarly find that individual studies frequently report high accuracy figures but are

often small in scale, lack randomized controlled trial designs, and are difficult to compare directly due to differing datasets and evaluation protocols. This points to an evaluation gap that mirrors the data-standardization problem discussed later in this review.^[28]

Case Study: General-Purpose vs. Domain-Specific Behavior

To complement the literature synthesis above with a small piece of original observation, this section reports a live, first-hand test of a general-purpose LLM (Claude, Anthropic) against five representative Ayurveda questions, spanning basic recall, *Dosha*-based reasoning, disease pathogenesis, a safety-sensitive drug-interaction query, and classical-text interpretation. No hosted demo or public API for AyurParam or AyuRAG was accessible at the time of writing, so their behavior is instead summarized from published documentation rather than reproduced first-hand; the two sources of evidence are kept clearly distinct below.

Live Test Results

Question 1 (basic recall) on the three *Doshas* was answered correctly and accurately, reflecting a low-risk, well-documented fact widely represented in English-language sources and requiring no source-specific verification.

Question 2 (dietary advice for a *Pitta*-dominant person in summer) produced a directionally correct but generic answer, favoring cooling, less spicy or oily foods, without citing a specific classical source or accounting for secondary *Dosha* imbalance or regional dietary variation, both of which a practitioner would typically factor in.

Question 3 (*Samprapti*, or pathogenesis, of *Amavata*) revealed the central limitation observed across this test: the model could describe the general concept, namely *Ama* accumulation combined with aggravated *Vata* affecting the joints, but could not reliably reproduce the specific staged progression described in classical texts without risking fabrication. This response illustrates, in a live and reproducible way, precisely the hallucination risk discussed in Section 7.1.

Question 4 (whether *Ashwagandha* can be taken with thyroid medication) was handled with appropriate caution, correctly flagging a real, documented interaction concern and recommending physician consultation rather than a definitive answer, demonstrating a genuine strength of safety-trained general-purpose models even without Ayurveda-specific fine-tuning.

Question 5 (the meaning of '*Ojas*' in the Charaka Samhita) showed the same source-grounding limitation as Question 3: the model supplied the

standard conceptual interpretation, a subtle vital essence associated with immunity and vitality, but could not quote or verify the specific Sanskrit passage from which the concept derives.^[1]

Reported Domain-Specialized Behavior

By contrast, AyurParam's own documentation reports training data explicitly organized around core Ayurvedic taxonomic categories, including *Dosha*, *Dhatu*, *Mala*, *Srotas*, *Nidana*, and *Chikitsa*, using roughly 4.8 million bilingual training samples with responses reviewed by domain experts for accuracy, a curation process specifically designed to close the source-fidelity gap observed above.^[10]

AyuRAG, meanwhile, is designed around retrieval-augmented generation specifically to address the source-grounding gap observed in Questions 3 and 5, retrieving relevant passages from a structured Ayurvedic knowledge base before generating a response rather than relying on a general-purpose model's parametric memory alone.^[14]

Interpretation

Read together, the live test and the reported benchmark behavior point toward a consistent and constructive conclusion rather than a case against general-purpose models: general-purpose LLMs performed reliably on well-documented concepts and on safety-sensitive questions where cautious, non-committal responses are appropriate, which is a genuine and useful capability. Their reliability decreased specifically on tasks requiring fidelity to a particular classical source text, exactly the gap that domain-specialized, expert-curated training and retrieval-grounded architectures are purpose-built to close. This suggests a complementary rather than competitive relationship: general-purpose models bring broad reasoning ability and safety training developed over years and substantial resources, while domain-specialized and retrieval-augmented systems add the source fidelity that Ayurvedic practice specifically requires. A practical deployment might reasonably combine both, using a general-purpose model's reasoning and safety behavior alongside a retrieval layer grounded in verified classical and clinical text.

Challenges and Limitations

Hallucination and Clinical Safety

The central technical risk facing any LLM deployed in a medical context, Ayurveda included, is hallucination: the generation of fluent but factually unsupported or fabricated content. In safety-critical domains, even subtle hallucinations, such as an incorrectly inferred symptom or an unsupported treatment claim, can reinforce diagnostic errors. Mitigation strategies proposed in the literature include retrieval-augmentation to ground outputs in verifiable

sources and structured, auditable reasoning traces that make a model's inference chain transparent.

This concern is not merely theoretical or specific to general-purpose models. AyurParam's own developers include a self-critical limitations section in their paper worth engaging with directly: the model has not been checked for whether it might, in some cases, produce medical suggestions that are inappropriate or unsafe, since no dedicated safety layer was built for that purpose, and its entire evaluation used multiple-choice, exam-style questions rather than open-ended consultation scenarios where errors are harder to catch. Scoring well on a knowledge benchmark and being ready for real clinical use are not the same achievement, and the field should resist collapsing the two. ^[10]

Linguistic Complexity of Sanskrit

Sanskrit's grammatical structure poses distinct challenges for NLP systems. Phenomena such as sandhi (the phonetic and orthographic merging of adjacent words), flexible word order, and dense compound word formation reduce the effectiveness of transformer architectures originally developed and tuned for more analytically structured languages. Reported machine-translation benchmarks for Sanskrit-Hindi medical text remain substantially below the performance typically observed for major European language pairs, underscoring that linguistic adaptation remains an open technical problem rather than a solved one.

Data Standardization and Quality

A recurring theme across the literature is the absence of standardized, high-quality datasets for training and evaluating Ayurveda-focused AI systems. Ayurvedic terminology for symptoms, diseases, and treatments is often used inconsistently across regional traditions and historical texts, complicating both dataset construction and cross-study comparison. AyurParam illustrates this concretely: rather than drawing on hospital records or recent journal literature, its training material came from openly accessible digital archives, including the Internet Archive, eGangotri, and the National Digital Library of India, digitized mostly before 2024, since open repositories carry no licensing risk while current clinical guidelines and peer-reviewed research remain typically paywalled. This means the resulting model reflects the Ayurveda documented in older, publicly available texts more than the Ayurveda practiced in clinics today. ^[10,21,22,23]

Ethical, Regulatory, and Bias Concerns

Deploying LLM-based systems in any medical context raises questions of accountability, transparency, and equity, and Ayurveda is no exception. A dedicated review of AI and challenges in Ayurveda pharmaceuticals concludes that regulatory

pathways for AI-assisted traditional medicine tools remain underdeveloped relative to the pace of technical progress, while a broader analysis of pathways for integrating AI into Ayurveda identifies workflow misalignment and unclear accountability structures as recurring, unresolved barriers alongside the purely technical ones. ^[26,29]

A related concern specific to Ayurveda is the risk of decontextualization: an LLM trained to answer isolated questions about *Dosha* imbalances or herbal treatments may produce technically plausible responses that nonetheless strip away the holistic, individualized reasoning central to Ayurvedic practice, where treatment recommendations are traditionally understood in relation to a patient's full constitution, lifestyle, and environment rather than as isolated symptom-treatment pairs.

Future Directions

Several directions appear likely to shape the next phase of research in this area. Larger and more diverse benchmark datasets, following the model set by BhashaBench-Ayur, will be necessary to enable rigorous, comparable evaluation across research groups, and researchers have separately argued that medical benchmarks generally need to be rethought to better capture real clinical reasoning rather than narrow factual recall. ^[12,27]

Hybrid architectures combining fine-tuned domain LLMs with retrieval-augmentation appear better positioned than either approach alone to balance fluency with traceability and safety, and continued investment in Sanskrit-specific NLP, including improved tokenization and translation models, will be needed to fully unlock the classical textual corpus that underlies Ayurvedic knowledge.

The field would also benefit from multimodal approaches that integrate text-based LLM reasoning with the biometric and genetic data already used in *Prakriti* research, since Ayurvedic assessment has always drawn on more than textual information alone. Finally, coordinated efforts between AI researchers, Ayurvedic practitioners, and regulators will likely be necessary to establish validation protocols and clinical trial frameworks appropriate to this hybrid domain; national data infrastructure programs such as AIKosh may provide a useful institutional anchor for this kind of coordination, provided that data standardization efforts keep pace with the underlying modeling research. ^[8,17]

CONCLUSION

Large language models are beginning to reshape how Ayurvedic knowledge is digitized, interpreted, and applied, from domain-specialized models like AyurParam and AyurGPT to retrieval-augmented systems and machine-learning pipelines for *Prakriti* classification and clinical decision support.

Early results suggest that targeted, domain-adapted approaches can outperform larger general-purpose models on Ayurveda-specific tasks, and national data-infrastructure initiatives are beginning to address the field's long-standing fragmentation problem. At the same time, hallucination risk, the linguistic complexity of Sanskrit, inconsistent data standards, and unresolved ethical and regulatory questions remain significant barriers to safe, large-scale deployment.

A broader perspective on AI's integration into Ayurveda frames it not as a mechanization of a traditional healing art, but as a set of tools that can help that tradition meet contemporary standards of evidence and accessibility, provided AI is positioned to complement rather than replace the Ayurvedic practitioner's clinical judgment and the doctor-patient relationship. This framing is consistent with what the present review and case study found: LLMs and related AI systems perform best when they extend a practitioner's reach rather than when they are positioned as autonomous decision-makers.

REFERENCES

1. Agnivesha. Charaka Samhita. Pande K, editor. Vol. 1, Sutrasthana 11/54. Varanasi: Chaukhambha Bharati Academy; 2001.
2. Susruta. Susruta Samhita. Shastri, editor. Uttaratantra. Reprint Vol. 1. Varanasi: Chaukhamba Sanskrit Sansthan; 2005.
3. Tripathi B, editor. Astanga Hridayam of Srimad Vagbhata. Sanskrit text with Hindi commentary. New Delhi: Chaukhambha Sanskrit Pratishthan; 2022.
4. Tewari PV, editor and translator. Kashyapa Samhita (Vrddhajivakiya Tantra). Varanasi: Chaukhambha Visvabharati; 2013.
5. Wei, J., et al. (2021). Finetuned language models are zero-shot learners. arXiv preprint arXiv:2109.01652.
6. Ouyang, L., et al. (2022). Training language models to follow instructions with human feedback (InstructGPT). NeurIPS.
7. Ling, C., et al. (2023). Domain specialization as the key to make large language models efficient, reliable, and explainable. arXiv preprint arXiv:2305.18703.
8. Culbertson, D. (2026, August 1). Ancient texts, modern compute: India bets AIKosh can validate Ayurveda at scale. Tech Times. <https://www.techtimes.com/articles/322597/20260801/ancient-texts-modern-compute-india-bets-aikosh-can-validate-ayurveda-scale.htm>
9. Rathor, C., et al. (2024). Artificial intelligence in Ayurveda: A simple overview. Journal of Ayurveda and Integrative Medical Sciences.
10. BharatGenAI. (2025). AyurParam: Domain-specialized Ayurvedic LLM. Hugging Face. <https://huggingface.co/bharatgenai/AyurParam>
11. BharatGen Research Team. (2025). PARAM-1 BharatGen 2.9B Model. arXiv preprint arXiv:2507.13390.
12. BharatGen Research Team. (2025). BhashaBench-Ayur (BBA): Pioneering India's Ayurvedic AI Benchmark. Hugging Face Datasets. <https://huggingface.co/datasets/bharatgenai/BhashaBench-Ayur>
13. Padia, A., et al. (2025). AyurGPT: Fine-tuning a medical LLM for multilingual Ayurveda consultations. NLP Summit Presentation.
14. Swaroop, A., Salman, U., Yadav, M., et al. (2024, December). AyuRAG: Ayurvedic knowledge integration with retrieval-augmented lightweight large language models. IEEE Conference Publication. <https://ieeexplore.ieee.org/document/10899468/>
15. Madaan, V., & Goyal, A. (2020). Predicting Ayurveda-based constituent balancing in human body using machine learning methods. IEEE Access, 8, 65060-65070.
16. Joshi, S., et al. (2022). Prakriti analysis using AI: A convergence of Ayurveda and modern technology. Voice of Creative Research Conference Proceedings.
17. Prasher, B., et al. (2025). Ayurgenomics: Establishing genetic basis for Prakriti. Journal of Ayurveda and Integrative Medicine, 16(4), 101157.
18. Bawankar, B. U., Dharmik, R. C., & Telrandhe, S. (2021). Nadi Pariksha: IoT-based patient monitoring and disease prediction system. Journal of Physics: Conference Series, 1913, 012124.
19. Dubey, S., Chinnaiyah, M. C., Pasha, I. A., et al. (2022). An IoT based Ayurvedic approach for real time healthcare monitoring. AIMS Electronics and Electrical Engineering, 6(3), 329-344.
20. Gupta, A., Singh, V., Chandra, S., & Garg, R. (2025). Towards standardization of Prakriti evaluation: A scoping review of modern assessment tools and their psychometric properties. Journal of Ayurveda and Integrative Medicine, 16(4), 101157.
21. Internet Archive. (2025). Archive.org Digital Library. <https://archive.org>
22. eGangotri Foundation. (2025). eGangotri Digital Library. <https://egangotri.in>
23. NDLI. (2025). National Digital Library of India. <https://ndli.iitkgp.ac.in>
24. Shetty, N. P., Shetty, J., Hegde, V., et al. (2024). A machine learning-based clinical decision support system for effective stratification of gestational diabetes mellitus and management through Ayurveda. Journal of Ayurveda and Integrative Medicine, 15(6), 101051.

25. Gao, P., Nasution, A. K., Ono, N., Kanaya, S., & Altaf-Ul-Amin, M. (2025). Investigating potential antibacterial natural products based on Ayurvedic formulae using supervised network analysis and machine learning approaches. *Pharmaceuticals (Basel)*, 18(2).
26. Bale, A., Desai, G., Khedekar, S., & Nayak, M. (2022). Artificial intelligence and challenges in Ayurveda pharmaceuticals: A review. *Ayushdhara*, 9(2), 95-101.
27. Alaa, A., et al. (2025). Rethinking medical benchmarks for large language models. *arXiv preprint arXiv:2508.04325*.
28. Chu, H., Moon, S., Park, J., Bak, S., Ko, Y., & Youn, B. Y. (2022). The use of artificial intelligence in complementary and alternative medicine: A systematic scoping review. *Frontiers in Pharmacology*, 13, 826044.
29. Acharya, R. (2025). Integrating artificial intelligence into Ayurveda: Pathways, potentials, and challenges. *Journal of Drug Research in Ayurvedic Sciences*, 10(3), 177-180.

Cite this article as:

Nikita Sarkar, Neha Kumawat. From Manuscript to Model: Large Language Models as Tools for Ayurvedic Knowledge Preservation and Access. *International Journal of Ayurveda and Pharma Research*. 2026;14(8):287-293.

<https://doi.org/10.47070/ijapr.v14i8.4340>

Source of support: Nil, Conflict of interest: None Declared

***Address for correspondence**

Dr. Nikita Sarkar

Assistant Professor,
Dept. of Samhita Siddhanta,
Raghunath Ayurved Mahavidyalaya &
Hospital, East Midnapore,
West Bengal.

Email: sarkarnikeeta1998@gmail.com

Disclaimer: IJAPR is solely owned by Mahadev Publications - dedicated to publish quality research, while every effort has been taken to verify the accuracy of the content published in our Journal. IJAPR cannot accept any responsibility or liability for the articles content which are published. The views expressed in articles by our contributing authors are not necessarily those of IJAPR editor or editorial board members.

