

Can AI replace Teachers? An Indian Knowledge Systems Perspective

Subhodeep Mukhopadhyay, Senior Research Fellow, Infinity Foundation India, Chennai, Tamil Nadu

Sudarshan Therani Nadathur, Director, Siddhanta Knowledge Foundation, Chennai, Tamil Nadu

Madhusudan Therani, Principal, Engkraft LLC, California, USA

Abstract

The question of whether artificial intelligence (AI) can replace teachers is an important one. Recent studies and public discourse increasingly suggest that AI systems can equal or even outperform human instructors in selected tasks such as content delivery, personalization, assessment, and feedback. This paper examines the question of AI replacing teachers from the perspective of Indian Knowledge Systems (IKS). The question we ask is whether AI can move beyond transmitting information and fulfill the deeper epistemic and civilizational role of the teacher. Among the different types of teachers in the Sanskrit tradition, we use the concept of *Apta*, a trustworthy authority on a topic, as seen in the *Nyaya* framework. We assess contemporary AI against the three criteria of an *Apta*: trustworthy source possessing direct knowledge, an intention to communicate truth, and the capacity to convey it properly. A desk review of current AI literature is conducted. Our results show that current AI systems do not qualify as an *Apta* since none of the three criteria are met. We conclude that while AI may simulate or partially perform certain instructional functions, it cannot be a complete substitute for the human teacher in the Indian philosophical sense.

Keywords: *Nyaya, Apta, Teacher, Indian Knowledge Systems, Generative AI, IKS*

Introduction

Can artificial intelligence replace the teacher? This question has now become central to discussions on the future of education. Rapid growth in intelligent tutoring systems and more recently, generative AI has steadily altered the educational landscape. Long before AI entered the classroom, online learning had already reduced the teacher's role in specific areas. Machines were significantly better than human teachers at certain tasks like standardized assessments and pre-recorded lecture delivery. Today sophisticated AI tools offer personalized tutoring with adaptive learning paths and real-time feedback (Gjermeni et al., 2026). Some tools even offer conversational guidance that feels like human interaction. Scholars therefore argue that the idea of teacher replacement is no longer merely rhetorical.

A growing body of literature has shown that AI-based tutoring systems can perform certain tasks exceptionally well. A recent experimental study on college students found that they learned "significantly more in less time" using an AI tutor. The students also felt 14 percent more engaged and 10 percent more motivated as compared to in-class active learning (Kestin et al., 2025). AI chat-based systems have also been shown to outperform human-only tutoring. A study by Google showed that students who received AI-based tutoring "were 5.5 percentage points more likely to solve novel problems on subsequent topics" than those who received human tutoring. There was

also qualitative evidence to suggest that the system “encouraged deeper reflection from students”(Team & Eedi, 2025).

However, critics argue that teaching is not just information delivery. A human teacher has empathy and moral discernment. He has emotional intelligence which allows him to exercise judgment in ways that AI cannot (Nikitina & Ishchenko, 2024). AI may supplement education, but not replace the teacher in any substantive sense, at least not in the foreseeable future. However, much of this debate, is framed within modern utilitarian and techno-utopian assumptions about education(Au et al., 2025; Segal, 1986).The teacher is implicitly treated as a content-deliverer or skill-transfer agent. Education is the medium which allows this transfer to happen. In an Indian context, such framings overlook the fact that a teacher has a far deeper role than being a mere instructional conduit.

Indian Knowledge Systems (IKS) offer a rich lens for exploring this question about whether AI can replace teachers. In the Hindu intellectual tradition, the teacher is not a singular figure. Terms such as *guru*, *acharya*, *upadhyaya*, *shikshaka*, *adhyapaka*, *pandit*, and *vidvan*, among others, point to distinct roles and responsibilities. Each of these deal with specific aspect of education and specific type of learning.A *guru*, for example, is not merely an instructor but a guide who transforms the student’s being. An *acharya* is one who teaches by both knowledge and conduct. A *shikshaka* or *adhyapaka* may be closer to the modern sense of educator. These are important distinctions because the answer to whether AI can replace the teacher depends also on which type of teacher is under discussion.

In this paper we ask: can AI replace the role of the teacher when the teacher is understood through an IKS perspective? To answer this, we use *Nyaya*, one of the classical schools of Hindu philosophy. *Nyaya* provides a sophisticated frame work for reasoning, debates, defining valid knowledge, and the means of knowing (Sarukkai, 2005). Among the means of knowledge (*pramanas*) recognized in *Nyaya*, *shabda*, or valid verbal testimony of a trustworthy person is an important one. Such a person is known as an *Apta* (Jha, 1984). He is someone who possesses direct knowledge of the matter, desires to communicate truthfully, and is capable of conveying that truth effectively to another. These three features together create the conditions under which testimony becomes a valid means of knowledge. The question we ask is therefore: can AI be regarded as an *Apta*?

Methodology

Our paper uses a desk review method. This is suitable because we address a conceptual question—whether AI can replace the teacher from an IKS perspective. We do so by examining and interpreting existing literature rather than collecting primary data. Our review is guided by the *nyaya* concept of *Apta* or trustworthy person who can give valid testimony, as the evaluation framework:

āptopadeśaḥ śabdaḥ | (NS 1.1.7) - The assertion of a reliable person is ‘Word.’

What this means is that the testimony of a reliable person can be considered knowledge. Such a person has three key attributes: “direct and right knowledge” of the topic, a commitment to communicating that truth, that is, a “desire to make know (to others) the things as he knows it”, and the capacity to convey truth properly and reliably(Jha, 1984, pp. 199–200). The desk review

examines current literature on AI against these attributes to assess whether AI can be regarded as a trustworthy teacher or *Apta*.

Results

Results are presented below for how well AI stacks up against: direct and right knowledge, commitment to communicating truth, and capacity to reliably convey truth.

Findings

Direct and Right Knowledge

The authority of an *apta* depends in part on genuine contact with and understanding of the object of knowledge. AI, however, does not know things through lived or causal encounter with the world. Large language models work by detecting and extending statistical patterns in text, not by directly apprehending reality (Bender et al., 2021). AI does not comprehend words or their meaning in the same way humans do. Any output produced by AI is created by sophisticated algorithms using complex mathematical calculations and predictive logic. One must not be mistaken in considering it genuine insight. The classic symbol grounding problem shows that a system may manipulate symbols successfully without any “understanding” of what those symbols refer to (Harnad, 1990). In the same way, AI also works with approximation representations of reality rather than reality as directly apprehended (Oka, 2025).

Often even this approximation is quite rudimentary, since AI systems have great conceptual difficulty in understanding context. The so-called frame problem in AI is an example of this. An AI system has no way of determining what is relevant and what can be ignored when something changes (Shanahan, 2016). Even when it can, the process is too complicated for an AI system. There is no notion of truth. In IKS, truth is not just having information; it is knowing what matters and why. AI systems can perform well in stable settings, mimicking human behavior, yet still fail when tacit assumptions or changing contexts become decisive.

Commitment to Communicating Truth

An *apta* is not only knowledgeable but also feels the imperative to communicate the truth to others. In the Hindu tradition, the drive to communicate one’s knowledge is the hallmark of a great Master. It is not because of specific motive like monetary gains or fame that the teacher wants to share his or her knowledge. Often in shastras he is shown as one driven by ideals like *karuna* and *daya* or compassion for the world. His purpose is welfare of the world as seen in Bhagavad Gita 3.20. He is the one who decides to bend down and raise “struggling individuals” (Swami Sivananda, 2025).

AI systems do not possess intention in the human sense. They have no self-purpose, and definitely nothing remotely resembling compassion driven urge to teach the world and communicate the truth. Their outputs are shaped by training data, optimization targets, safety tuning, moderation rules, and product-level policy constraints. This means they may refuse some questions or redirect some lines of inquiry. Sometimes they engage unevenly with competing viewpoints, and prioritize one point of view. Policies primarily decide what AI can communicate. A further problem is bias. Current research strongly shows that LLMs can reproduce, and sometimes amplify social, political, cultural, and religious biases present in their training data. There is extensive evidence that LLMs encode

stereotypes and patterned unfairness even in embeddings and probabilities, which in turn shape the generated text (Caliskan et al., 2017; Huang et al., 2023; Jackson, 2021; Mukhopadhyay & Reddy, 2023).

Thus, there is no commitment to communicate truth in any form. Normative policies and prevailing worldviews shape what AIs produce.

Capacity to Reliably Convey Truth

LLMs are designed primarily to predict plausible next tokens. Their purpose is not to verify truth. Their architecture is optimized for linguistic fluency and coherence, and not for a built-in correspondence relation with reality. This is why a response may sound authoritative and internally consistent while still being false, sometimes known as “stochastic parrots” (Bender et al., 2021). This weakness becomes most visible in the problem of hallucination. Current research and risk frameworks recognize that generative AI can produce fluent but false, or fabricated content. Such hallucination is not a rare accident but a structural issue tied to training, and prediction (Sun et al., 2024). A source that can confidently generate plausible falsehoods cannot qualify as an *Apta*.

Summary and Implications

AI does not meet even a single defining characteristic of an *Apta*. It lacks direct knowledge of the object, it has no intrinsic commitment to truth, and it has no self-aware capacity to convey knowledge.

Direct apprehension, is not equivalent to merely recombining previously seen verbal material as AI does. AI does not “perceive” the world as a conscious knower does. Its relation to reality is mediated through data traces produced by others. Hence, even when its output is useful, its “knowledge” is derivative and indirect. AI systems do not meet the test of commitment to spread the truth. Unlike an *apta*, AI does not possess a self-grounded desire to speak truth as truth. It follows optimization and policy. Moreover, if the system’s outputs are shaped by opaque sampling biases, or political and social skew, then there cannot be any obligation to spread right knowledge. While some people may find some use, by and large it remains prejudicial and unreliable. Valid testimony depends on a source whose speech is truth-bearing because it arises from knowledge. In LLMs, by contrast, truth is not the core organizing principle of the architecture. The system can incidentally approximate truth from time and time, but it is not inherently ordered toward truth as such. The clinching factor, of course, is the irreconcilably flawed phenomenon of hallucination. A source that confidently generates falsehoods while presenting them in the form of knowledge cannot qualify as a trustworthy verbal authority.

Analysis

If AI is not an *apta*, then what it produces cannot qualify as *aptopadesha* or valid *shabda*. And if it is not valid *shabda*, it cannot function as *pramana*, that is, as a valid means of knowledge. What AI produces is a simply pattern-based response or probabilistic text generation. It may of course be useful for many different categories of tasks and different groups of people may find value in what AI produces. But it is not knowledge-bearing testimony.

This limitation becomes even clearer when viewed through the broader distinction between *sthula* and *sukshma* realms. In Vedic knowledge systems the universe and all its objects including operates at various levels: *sthula* (gross), *sukshma* (subtle) and *karana* (causal). Knowledge does not operate only at the material or *sthula* level but across all three realms. AI however operates primarily at the level of the *sthula* or the material world. Even here its grasp is only partial. There is no symbol grounding in AI, that is, AI cannot recognize the real-world object referred to by a word. AI cannot relate the word “cat” to the object “cat”, despite the fact that any GenAI system can easily generate a fluent and reasonably accurate essay on cat. In that sense, even its relation to the *sthula* world is indirect and ungrounded.

An *apta* is trusted not merely because he repeats correct statements. But because he knows in a fuller sense. He is grounded in knowledge and can communicate from that realization. AI has no access to such depth. For this reason, AI may at best simulate fragments of the gross domain, but it cannot serve as a knower in the deeper sense required for valid testimony. This is especially important since IKS is premised on the *guru-shishya parampara* (teacher-student lineage). This is not mere transfer of propositional content. It is a faculty which facilitates *aparoksha-gyana* or direct experience. This knowledge is qualitatively different from *paroksha-gyana* or mediated, second-hand knowledge.

Furthermore, the *apta* must possess *viveka*. This is discriminative wisdom that distinguishes the real from the apparent, the essential from the incidental, the truth from falsehood, and *dharma* from *adharma*. *Viveka* is not an algorithm. It is a function of conscious discernment ripened through experience and *tapas* (disciplined practice). An AI system lacks sentience and hence cannot exercise *viveka* in any meaningful sense. It can, of course, produce outputs that mimic discriminative judgment. But such mimicry are but simply appearance without substance.

Conclusion

The question of whether AI can assume the role of a teacher is not merely technological; it is, at its core, a philosophical question about the nature of knowledge and meaning. Yet as we have seen AI will not be able to replace a teacher in the IKS sense. With every passing year, we may have better, more optimized and far more powerful AI system, but they will still operate at the material level. An AI will not be able to replace an *Apta*.

The larger question is whether AI will be able to replace other types of teachers as well within IKS. As noted earlier, there are different types of teachers in the Hindu tradition (Table 1). This could be an interesting area of research for future studies. But let us look at an example. One may think the role of *adhypaka* or *shikshaka* may get partially or wholly replaced since AI excels at the instructional level tasks. However, deeper research is required since the way *kala* and *shiksha* is conceptualized in IKS is quite distinct from “arts/skills” and “education” in the modern Western knowledge system sense.

Table 1 *Types of Teachers in the Hindu Tradition*

Role	Category	Subject Mastery
<i>Guru</i>	Spiritual	Mastery of the Self and <i>Shastras</i> .
<i>Acharya</i>	Master	Complete mastery of the <i>Vedas</i> and <i>Vedangas</i> .
<i>Apta</i>	Authority	Direct knowledge of a subject.
<i>Upadhyaya</i>	Professional (Vedic)	Expertise in a specific branch of the <i>Vedas/Vedangas</i> .
<i>Pandit</i>	Intellectual	Expertise in specific <i>Shastras</i> (Logic, Grammar etc.).
<i>Vidvan</i>	Learned	Broad knowledge across multiple fields.
<i>Adhyapaka</i>	Instructional	General academic subjects.
<i>Shikshaka</i>	Skill-based	Specific subjects, crafts, or technical skills.

This does not mean AI has no pedagogical utility. As an *upaya* or an instrumental means, AI can serve important functions: organizing information, summarizing, ideating, automating, and so on. It helps make *paroksha-gyana* widely available. These are not trivial contributions. But they are contributions of a sophisticated tool, not of a teacher. The mirror can reflect light, but it is not the sun. The *Apta* must be one who has directly known. Until AI can satisfy this condition, which would require it to possess *chetana* or awareness, the *guru-shishya parampara* remains a distinctly and irreducibly human institution. The *Apta* is not going anywhere.

References

1. Au, L. N., Do, M. T. N., & Nguyen, H. D. P. (2025). Controversies surrounding utilitarian learners and their learning behaviours: Do educators know the facts? *Creative Education*, 16(3), 221–229. <https://doi.org/10.4236/ce.2025.163013>
2. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, 610–623. <https://doi.org/10.1145/3442188.3445922>
3. Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186. <https://doi.org/10.1126/science.aal4230>
4. Gjermeni, F., Sena, I., Prodani, F., & Pereira, A. I. (2026). AI-powered tutoring for personalized learning. In A. I. Pereira, F. P. Fernandes, J. P. Coelho, J. P. Teixeira, J. Lima, M. F. Pacheco, L. Oneto, & R. P. Lopes (Eds.), *Optimization, learning algorithms and applications* (Vol. 2618, pp. 191–204). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-00140-5_13
5. Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1–3), 335–346. [https://doi.org/10.1016/0167-2789\(90\)90087-6](https://doi.org/10.1016/0167-2789(90)90087-6)
6. Huang, C., Zhang, Z., Mao, B., & Yao, X. (2023). An overview of artificial intelligence ethics. *IEEE Transactions on Artificial Intelligence*, 4(4), 799–819. <https://doi.org/10.1109/TAI.2022.3194503>

7. Jackson, M. (2021). Artificial intelligence & algorithmic bias: The issues with technology reflecting history & humans. *Journal of Business & Technology Law*, 16(2), 299.
8. Jha, G. (1984). *The Nyāya-sūtras of Gauṭama: With the Bhāṣya of Vātsyāyana and the Vārtika of Uddyotakara*. Motilal Banarsidass.
9. Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15(1), 17458.
<https://doi.org/10.1038/s41598-025-97652-6>
10. Mukhopadhyay, S., & Reddy, D. (2023). Artificial intelligence as an enabler of Western universalism. In R. Malhotra, T. N. Sudarshan, & M. Sastry (Eds.), *The power of future machines: Essays on artificial intelligence* (pp. 241–263). BluOne Ink LLP.
11. Nikitina, I., & Ishchenko, T. (2024). The impact of AI on teachers: Support or replacement? *Scientific Journal of Polonia University*, 65(4), 93–99. <https://doi.org/10.23856/6511>
12. Oka, S. (2025). *Evaluating large language models on the frame and symbol grounding problems: A zero-shot benchmark* (Version 1). arXiv.
<https://doi.org/10.48550/ARXIV.2506.07896>
13. Sarukkai, S. (2005). *Indian philosophy and philosophy of science* (1st ed.). Centre for Studies in Civilizations.
14. Segal, H. P. (1986). The technological utopians. In J. J. Corn (Ed.), *Imagining tomorrow: History, technology, and the American future*. MIT Press.
15. Shanahan, M. (2016). The frame problem. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2016 ed.). Metaphysics Research Lab, Stanford University.
<https://plato.stanford.edu/archives/spr2016/entries/frame-problem/>
16. Sun, Y., Sheng, D., Zhou, Z., & Wu, Y. (2024). AI hallucination: Towards a comprehensive classification of distorted information in artificial intelligence-generated content. *Humanities and Social Sciences Communications*, 11(1), 1278.
<https://doi.org/10.1057/s41599-024-03811-x>
17. Swami Sivananda. (2025, July 6). The role of the guru. *The Divine Life Society*.
<https://www.sivanandaonline.org/newsupdates/the-role-of-the-guru/>
18. Team, L., & Eedi. (2025). *AI tutoring can safely and effectively support students: An exploratory RCT in UK classrooms* (arXiv:2512.23633). arXiv.
<https://doi.org/10.48550/arXiv.2512.23633>