The ‘Sarvam’ [Everything] in India’s AI Strategy

April 10, 2026



Two months after SarvamAI unveiled its family of ‘sovereign models’ at the AI Impact Summit—with PM Modi on stage—India’s foremost AI startup is reportedly in talks to raise funds between $300 to $350 million at a valuation of $1.5 billion. The funding round is set to close as early as next week.

SarvamAI is arguably the most strategically significant AI company in India at present. As such, the state of its finances carries implications that extend far beyond the venture capital ecosystem, potentially shaping the trajectory of India’s ‘sovereign’ AI sector. It is therefore pertinent to critically assess the evolution of Sarvam’s role in India’s AI ecosystem, its relationship with the GoI, and its broader mission. Doing so will yield insights not only into India’s approach to AI, but also into the future of AI development in the country.

The foreign AI problem

The reason SarvamAI is so strategically vital to India is

straightforward: almost all of the roughly 900 GenAI startups in the country have built applications and services on top of licensed foundational models developed by Big Tech firms abroad, or on open-source models. Sarvam is perhaps the sole prominent exception—almost by mandate (as discussed later).

There are several reasons for India's dependence on 'foreign' AI, but foremost is a lack of capital. Last year, for instance, Indian startups raised just over $1.3 billion, amounting to a meagre 0.6% share of the global AI pie. For further context, OpenAI alone raised $122 billion in the last financial year, even as it remained unprofitable.

India's first AI strategy

Despite this reality, or perhaps because of it, India's approach to AI—from the period between the release of ChatGPT in 2022 through early 2025—was largely informed by narratives around the country's history of 'frugal innovation'. The belief was that Indian AI startups could develop low-cost products tailored to the domestic market (LLMs in Indian languages, and so on) and sell them at scale.

As Infosys co-founder and 'the man behind Aadhaar,' Nandan Nilekani, declared in October 2024, "Our goal should not be to build one more LLM. Let the big boys in the [Silicon] Valley do it, and spend billions of dollars." In other words, Nilekani argued that India's competitive advantage lay in the creation of "use cases."

At the time, the logic was solid (and in some ways, it still is). For instance, SarvamAI had made a name for itself through voice AI products, built on open-source models trained on Indian languages. Others had built AI enterprise solutions with Indic scripts, and so on. Meanwhile, India's AI adoption rates were (and still are) far higher than the global average.

The Sarvam-led pivot and the GoI's hand

Then, in January 2025, the release of China's DeepSeek disrupted both the global stock market and India's assumptions about AI development. As one prominent Indian AI analyst noted, "Deepseek laid bare the truth: India wasn't in the frugal-innovation game at all, and China was far ahead of everyone." By this time, it was further evident that Indian AI startups' products—built on largely self-professed advantages in Indic languages use cases—were neither cheap nor capable of matching the benchmark performance of products developed in the US.[1]

In this context, just weeks after DeepSeek's launch, the GoI declared that the IndianAI Mission would build 'sovereign AI' (DIPTEL #111). This was a major and expansive shift away from the program's primary strategy since inception: to reduce the cost of AI innovation through subsidized compute power for Indian startups and research labs. The idea now was to build foundational models at home, which other companies could utilize to build products tailored for the Indian market and beyond. The process would not only give India an edge in the global AI race but also deliver a level of much-needed strategic tech independence.

From this point onward, SarvamAI assumes enormous significance. For reasons that remain unclear, the GoI handpicked the startup to lead the development of India's first foundational model. As one AI founder stated at the time, "There's no clarity on why the company was selected… It's not possible to tell if they just showed a fancy demo or if due diligence was actually done. There's no transparency."

Sarvam was a pre-revenue company, founded just two years ago in 2023, making it an odd choice for a government tender, and although it had impressive tech capabilities relative to its size, it was far from the only advanced AI lab in the country. Experts argued that India must back multiple competitors to hedge its bets, but the advice was unheeded, and Sarvam

emerged as the chosen one.

A lock-in effect

Meanwhile, critics stated that Sarvam's exclusive arrangement with the government could create a 'lock-in' effect, with its AI products likely deployed across the country's digital public infrastructure to automate government services in the future. That process, to some extent, is already underway. Sarvam's technology has been deployed in India's biometric identification system (Aadhaar), the Indian central banks' life insurance products, and other critical areas. Through this lens, as a prominent AI expert observed, "Sarvam AI resembles an arm of the Indian government's AI strategy more than a conventional VC-backed startup."

Earlier this year, the GoI appeared to partially recalibrate its approach by including 12 other players, such as BharatGen (incubated in IIT Bombay), in IndiaAI Mission's next round of sovereign AI development. But unsurprisingly, Sarvam appears far ahead at present. Its aforementioned models released in February—the Sarvam 105B and the smaller Sarvam 30B—were trained across 22 Indic languages on an entirely domestic tech stack. Benchmarks suggest strong performance overall (relative to the cost of its development), and on certain Indic-specific tasks, the company claims performance that surpass far larger models such as GPT or Gemini.

Sarvam achieved all this with roughly $50 million raised through its Series A rounds, and another $25 million in government support. At first glance, then, news that SarvamAI is set to raise $300-$350 million appears to mark a transformative moment in India's AI journey.

However, this development must be read with two important caveats:

The benchmark problem

First, nearly all of Sarvam's reported benchmark performance remains self-reported and has yet to be independently verified. Presently, the only Sarvam AI model to appear on the Hugging Face Open LLM Leaderboard—a prominent, community-driven platform that evaluates and ranks models across multiple standardized benchmarks—is OpenHathi, an open-source LLM built two years ago, which ranks beyond 4000 globally. Its new models also do not have an Arena rank (a prominent evaluation platform for LLMs based on human preferences), and even more importantly, Sarvam has not yet released any formal arXiv papers (with methodology, ablation studies, and peer-review) for either 105B or 30B models.

Meanwhile, IndiVibe, the prominent Indic language evaluation benchmark, was designed by Sarvam itself, raising obvious concerns about conflicts of interest. A few other benchmarks, such as MILU or IndicGLUE, were created by AI4Bharat, whose founders have since joined Sarvam.

This does not necessarily imply bad faith; the problem is, in large part, structural. As a prominent AI expert recently explained,

"India has begun to build evaluation efforts for Indic AI, but it still lacks an independent, nationally trusted scoreboard with the reach, rigour and legitimacy to arbitrate ambitious claims like Sarvam's. There is no Indian equivalent of Arena in terms of visibility and community trust. There is no HELM-scale evaluation harness maintained by an institution independent of the model builders."

This is not merely an academic problem. If Sarvam's models underperform on claimed benchmarks—particularly in Indic use cases—they could directly affect service delivery in domains that 'sovereign AI' is meant to strengthen, such as digital governance. In this context, the GoI will likely accelerate the development of AI regulatory infrastructure (which includes benchmark performance evaluation), likely under the

IndiaAI Mission's Safety & Trusted AI pillar.

As indigenous AI products proliferate, and especially if the GoI subsidizes their development, the need for independent evaluation will only intensify. Moreover, future investments in India's AI sector will also likely be bolstered by, if not contingent upon, the same.

What is the money for?

If Sarvam's models remain insufficiently evaluated, how is the company poised to raise nearly a quarter of the total capital raised by India's AI startup ecosystem last year?

Part of the answer may lie in visibility gains following the AI Impact Summit and sustained government backing. However, there is another possibility—one that complicates the narrative of strategic gain.

In January, SarvamAI and the Tamil Nadu state government signed an MoU to build 'India's first Sovereign AI park' at a cost of over $1 billion, which includes data centers that exclusively host government data. At that time, when Sarvam had raised just about the aforementioned $50 million, a state minister had declared, "Sarvam will bring in the investment for the project while the government will provide land and necessary support."

It is therefore plausible that a significant portion of Sarvam's incoming capital will be directed not toward AI research or model development, but toward physical infrastructure. If so, the strategic implications for India's sovereign AI ambitions may be far more limited than they appear, as Sarvam's ability to convert its early momentum in frontier AI into a durable national advantage becomes uncertain.

[1] In early 2025 onwards, several AI firms in the US launched cheap AI products, such as Google's Gemini Flash series, built on smaller but efficient models, that targetted markets like India. For many automated tasks at that time, even in local use cases such as translation in Indic languages, these products eroded the price advantage (per million tokens) that Indian AI startups (like Sarvam) assumed they would have.