

India's AI Strategic Paralysis: The 69-Month Freeze That Cost a Generation

A policy analysis of how long India took to respond to the foundation-model shift in AI, measured on dated policy documents against China, the United States and the EU, with evidence on research, talent, compute and models through 2026.

Saket Poswal

Author

Oct 14, 2025 • Updated: Sep 29, 2026

ABSTRACT

Between June 2018, when NITI Aayog published India's National Strategy for Artificial Intelligence, and March 2024, when the Union Cabinet approved the IndiaAI Mission, India did not revise its national AI strategy or fund national AI compute, even as the field shifted from task-specific models to large foundation models. This paper asks how long India took to respond to that shift, why, and what the delay cost. It introduces a simple, auditable measure, policy adaptation lag: the interval between a public paradigm-shift signal (the release of GPT-3 in June 2020 and of ChatGPT in November 2022) and a funded or binding national response, measured on dated primary documents. On this measure India's funded response came about 45 months after GPT-3 and 16 months after ChatGPT, and its first publicly launched sovereign foundation models came about 39 months after ChatGPT, in February 2026. The paper sets these intervals beside the responses of China, the United States and the European Union, noting that those responses differ in kind. It then reviews capability indicators: a 1.4% share of papers at top AI conferences (2018–2023), the world's highest AI skill penetration alongside the largest net outflow of AI talent in 2025, compute that grew to more than 38,000 subsidised GPUs by December 2025, and models from IndicTrans2 to Sarvam-105B. The period was not a complete freeze: language-technology programmes such as Bhashini and AI4Bharat were genuine investments in the model layer. The paper proposes five hypothesised causes of the lag, weighs counter-arguments including the case for a fast-follower strategy, and recommends institutional mechanisms, such as scheduled strategy reviews and paradigm-shift triggers, designed to shorten future lags.

Author

Saket Poswal · [ORCID 0009-0009-8574-9953](#) · [Google Scholar](#)

Type	Analysis (working paper, not peer reviewed)
Published	14 October 2025
Last revised	29 September 2026
Method	Policy analysis using dated primary documents; construction of a policy adaptation lag measure; comparative analysis of national AI policy timelines; secondary analysis of published indicators on research output, talent, compute and models
License	CC BY 4.0
DOI	10.5281/zenodo.23011034
Keywords	India AI strategy · IndiaAI Mission · policy adaptation lag · foundation models · sovereign AI · AI talent migration · AI compute · NITI Aayog · Bhashini · AI4Bharat · Sarvam · BharatGen · technology policy

Working paper · Version 2.0 (29 September 2026)

Revision note. Version 2.0 substantially revises the October 2025 paper. Its central argument, that India responded late to the foundation-model shift, is retained but now measured on dated primary documents using an explicit lag measure rather than asserted. The paper has been updated with developments through 2026, including the first IndiaAI foundation-model launches, which the earlier version's claim of "zero models" preceded. Adaptation-cycle figures for China, the United States and the EU that were not traceable to sources were removed and replaced with dated policy milestones. Unverifiable course-by-course curriculum claims and news citations that pointed only to publishers' home pages were removed. Counter-evidence (Bhashini, AI4Bharat and the fast-follower argument) has been added, along with a limitations section. The "90-day emergency reset" has been replaced with durable institutional recommendations. The title's "84-month" figure was an arithmetic error; June 2018 to March 2024 is 69 months, and the title has been corrected.

1. Introduction

In December 2024, the Chinese laboratory DeepSeek published the technical report for DeepSeek-V3, a large language model competitive with leading Western systems. The report put the cost of its final training run at about US\$5.6 million in GPU time, excluding prior research and experiments (DeepSeek-AI, 2024). For countries that had assumed frontier AI required resources on the scale of the largest U.S. firms, the report was a signal that the entry barrier was lower than believed and that lost time, not lack of money, might be the binding constraint.

For India, the signal raised an uncomfortable question. India published a thoughtful national AI strategy in 2018 (NITI Aayog, 2018) and has one of the world's largest AI talent pools. Yet its first publicly launched sovereign foundation models under a national programme appeared only in February 2026. How long did India take to respond to the foundation-model shift, why, and what did the delay cost?

This paper makes four contributions:

1. A measure. *Policy adaptation lag* is the interval between a public signal that the technological paradigm has shifted and a funded or binding national response, measured on dated primary documents so that others can check and update it (Section 3).
 2. A comparative timeline for India, China, the United States and the European Union, with the differences in kind between their responses made explicit (Section 3).
 3. An evidence review of capability indicators: research output, talent, compute and models (Section 4).
 4. A causal account and recommendations, with the causes framed as testable hypotheses and the recommendations as institutional mechanisms rather than one-off actions (Sections 5–8).
-

2. Background

2.1 The 2018 strategy

NITI Aayog's *National Strategy for Artificial Intelligence* (2018), branded #AIforAll, identified five priority sectors (healthcare, agriculture, education, smart cities and mobility), proposed research institutions, and emphasised AI for inclusive development. Its orientation was toward *applying* AI to Indian problems. That was a reasonable choice at the time: the dominant paradigm was task-specific models trained for particular applications, and the strategy's sectoral focus fitted it.

2.2 The paradigm shift

The strategy appeared a year after the transformer architecture was introduced (Vaswani et al., 2017) and just before a rapid change in what AI systems were and who could build them. In June 2020, GPT-3 showed that a single very large pretrained model could perform many tasks from a few examples, without task-specific training (Brown et al., 2020). In November 2022, ChatGPT brought that capability to the public. From that point, the strategic question for any country was no longer only how to *apply* AI but whether it had access to the *foundation layer* (large models, the compute to train them, and the data and talent to build them) on which applications increasingly depended.

2.3 India between 2018 and 2024: not a complete freeze

The title of this paper refers to the interval between the 2018 strategy and the 2024 IndiaAI Mission during which the national strategy was not revised and no national AI compute was funded. That interval ran from

June 2018 to March 2024, about 69 months. The first version’s title gave this interval as 84 months; that was an arithmetic error, corrected here. Several things did happen during the interval:

- AIRAWAT (2020). NITI Aayog proposed an AI-specific cloud computing infrastructure, benchmarked India’s compute against other countries and outlined a governance model (NITI Aayog, 2020). It was an approach paper; the national compute it envisaged was funded only in 2024.
- Responsible AI (2021). NITI Aayog published principles for responsible AI (NITI Aayog, 2021).
- Bhashini (2022). The Ministry of Electronics and Information Technology launched the National Language Translation Mission in July 2022 to provide language technology in 22 Indian languages as a public resource.
- AI4Bharat. Researchers at IIT Madras released open models for Indian languages, including IndicTrans2, the first open-source transformer translation system covering all 22 scheduled languages (Gala et al., 2023).

These were real investments in Indian-language models. What was missing was a national response to the *general-purpose* foundation-model layer, the compute to train such models and a strategy revised to reflect the new paradigm.

3. Measuring Policy Adaptation Lag

3.1 Definition

For a country c and a paradigm-shift signal s with public date d_s , the policy adaptation lag is $L_{c,s} = d_r - d_s$, where d_r is the date of the first funded or binding national response addressing the new paradigm, established from a dated primary document. Announcements, consultations and approach papers without funding or legal force do not count as d_r , but are recorded. Two signals are used: GPT-3 (June 2020), the signal to specialists, and ChatGPT (November 2022), the signal to everyone.

3.2 Timeline

Table 1. Dated national responses to the foundation-model shift

Country	Milestone	Date	Type
India	National Strategy for AI (NITI Aayog)	Jun 2018	Strategy, pre-shift
	AIRAWAT approach paper	Jan 2020	Proposal (unfunded)
	Bhashini launched	Jul 2022	Funded language-technology programme
	IndiaAI Mission approved, ₹10,371.92 crore	7 Mar 2024	Funded (compute, foundation models, data, skills)
	Foundation-model developers selected	Apr 2025	Implementation

Country	Milestone	Date	Type
	AI Governance Guidelines	Nov 2025	Non-binding governance framework
	First sovereign foundation models launched (Sarvam-30B, Sarvam-105B and others)	Feb 2026	Output
China	New Generation AI Development Plan	Jul 2017	Strategy
	Interim Measures for Generative AI Services	Jul 2023 (in force Aug 2023)	Binding regulation
United States	Executive Order 14110 on safe, secure and trustworthy AI	Oct 2023	Binding on federal agencies
	EO 14110 revoked; <i>America's AI Action Plan</i>	Jan 2025; Jul 2025	Policy reversal; strategy
European Union	AI Act proposed by the Commission	Apr 2021	Legislative proposal
	AI Act adopted with general-purpose AI provisions, in force	Aug 2024	Binding regulation

Sources: NITI Aayog (2018, 2020); Press Information Bureau (2024, 2025); Ministry of Electronics and Information Technology (2025); Business Standard (2026); State Council (2017); Cyberspace Administration of China et al. (2023); Executive Office of the President (2023); The White House (2025); European Union (2024).

3.3 Results

- India: about 45 months from GPT-3, and 16 months from ChatGPT, to a funded response (the IndiaAI Mission); about 39 months from ChatGPT to the public launch of foundation models under that programme.
- China: about 8 months from ChatGPT to binding generative-AI regulation.
- United States: about 11 months from ChatGPT to an executive order.
- European Union: general-purpose AI provisions were added to legislation already in progress, entering into force about 21 months after ChatGPT.

3.4 Interpreting the comparison

The comparison needs care. China's, the United States' and the EU's fastest responses were *regulatory*, while India's was *industrial*: funding compute and models. Regulatory responses can be issued faster than industrial ones can be built, and China and the United States already had large private foundation-model efforts that did not depend on state funding. The fairest reading is not that India was uniquely slow on every dimension. Two findings stand out: India had no funded national response to the foundation layer until 2024, about four years after the paradigm shift was visible to specialists; and India's national strategy document was not revised during that period, so the shift was never formally incorporated into national strategy before the Mission.

4. Evidence: What the Lag Coincided With

The indicators below show conditions during and after the lag. They are associations, not proof that the lag caused them.

4.1 Research output

An analysis by the AI accelerator Change Engine, reported by The Wire, found that India contributed 1.4% of papers at the top ten AI conferences from 2018 to 2023, ranking 14th, compared with 30.4% for the United States and 22.8% for China. India's number of papers grew at 15.5% a year over 2014–2023, compared with 20–30% for the leading countries (Change Engine, 2024). This is an industry analysis rather than a peer-reviewed bibliometric study, and conference share is only one measure of research strength, but the gap it shows is large.

4.2 Talent: world-leading skills, world-leading outflow

The *AI Index 2026* reports that India had the world's highest relative AI skill penetration on LinkedIn, with AI skills appearing in profiles at 3.0 times the global average, and the second-largest pool of AI researchers and inventors (50,460 in 2025, after the United States). It also reports that India recorded the largest net outflow of AI talent of any country in the dataset in 2025 (Stanford Institute for Human-Centered Artificial Intelligence, 2026). The combination is the core of the talent problem. India produces AI skills at scale but, during the lag, offered few domestic opportunities to work at the frontier.

4.3 Compute

Once funded, compute grew quickly. The IndiaAI Mission's original target was about 10,000 GPUs. By December 2025, more than 38,000 GPUs had been onboarded through a national compute portal and made available to start-ups, researchers and academic institutions at subsidised rates of about ₹65 per GPU-hour (Press Information Bureau, 2024, 2025). That this scale could be reached within about 21 months of approval suggests that compute was a constraint of *decision*, not of capacity, during the lag.

4.4 Models

India's model-building record shows the same pattern of capability arriving once funding and direction arrived:

- 2023: IndicTrans2, open translation models for all 22 scheduled languages (Gala et al., 2023).
- 2025: PARAM-1, a 2.9-billion-parameter Hindi–English model trained from scratch by the government-supported BharatGen consortium (BharatGen, 2025).
- February 2026: Sarvam-30B and Sarvam-105B (a mixture-of-experts model with about 10 billion active parameters), trained from scratch on IndiaAI compute, supporting 22 Indian languages and

later released as open weights (Business Standard, 2026).

4.5 What the DeepSeek comparison does and does not show

DeepSeek-V3's reported US\$5.6 million covers only the final training run, at an assumed rental price for GPUs, and excludes the research, experiments, staff and infrastructure behind it (DeepSeek-AI, 2024). It does not show that a frontier model can be built for US\$5.6 million from a standing start. It does show that the marginal cost of training a competitive model had fallen to a level well within the reach of a national programme, and that the main barriers were accumulated expertise and early commitment, precisely what a lag erodes.

5. Why the Lag? Five Hypotheses

The first version of this paper presented five root causes as established. The evidence available supports them only as hypotheses; each is stated with what would test it.

H1. No scheduled review. The 2018 strategy had no built-in revision cycle and no designated trigger for revisiting it, so a paradigm shift did not automatically prompt a strategic response. *Test:* compare lag across countries or sectors with and without mandated review cycles.

H2. Academic incentives. Faculty careers reward publication in established subfields, and training frontier models requires compute that Indian universities largely lacked before 2024, so academic priorities could not shift quickly. *Test:* track changes in the topical composition of Indian AI publications before and after the 2024 compute expansion.

H3. Application-layer success. India's success with digital public infrastructure (identity, payments and data exchange) oriented policy toward applications and platforms, a strength that may have made the foundation layer seem someone else's business. *Test:* analyse the content of AI policy documents and budget speeches for application versus infrastructure framing over time.

H4. Funding rigidity. Multi-year budget and approval cycles are slow relative to a field in which capabilities changed substantially within a year. *Test:* measure time from proposal to approval for AI programmes relative to other science programmes.

H5. Weak error-correction mechanisms. Public institutions have few routine, low-cost ways to say that a strategy has been overtaken by events. The first version framed this as a cultural failing ("pride over pragmatism"); an institutional description is more accurate and more actionable. *Test:* examine whether the lag shortened after the 2024 Mission established a standing implementing body.

6. Counter-Arguments

“Application-first was rational.” For a country with limited compute budgets and urgent development needs, applying existing AI to health, agriculture and public services may have yielded more benefit per rupee than training foundation models. *Response*: partly right. But reliance on foreign foundation models carries risks for language coverage, cost and strategic autonomy, and the talent outflow in Section 4.2 suggests that absence from the foundation layer also affected the application layer by driving away the people who build both.

“Fast-follower strategy.” With open-weight models widely available, a country can adopt and adapt frontier models without training its own, so a lag costs little. *Response*: open models reduced the cost of the lag, and much Indian work sensibly builds on them. But fast-following requires the ability to train and adapt models, which depends on compute and talent built during the lag years, not after.

“India was not idle.” Bhashini and AI4Bharat were investments in the model layer for Indian languages (Section 2.3). *Response*: accepted. This is why the paper measures a lag in the *general-purpose foundation layer and national strategy*, not a total absence of activity.

“The lag is over.” By 2026, India had funded compute at scale and launched sovereign models. *Response*: the lag is over, but its costs, in talent outflow and research share, persist, and without institutional changes the next paradigm shift may produce another.

7. The 2024–2026 Response and Remaining Gaps

The IndiaAI Mission, approved on 7 March 2024 with an outlay of ₹10,371.92 crore over five years, funds compute capacity, an innovation centre for indigenous foundation models, a datasets platform, application development, skills, start-up financing and safe and trusted AI (Press Information Bureau, 2024). By the measures available in 2026, its implementation has been fast: compute exceeded its target within about 21 months, and foundation models followed within about ten months of developer selection.

Gaps remain:

- Research share. Conference share responds slowly to funding and remains low (Section 4.1).
 - Talent retention. The 2025 net outflow is the largest recorded (Section 4.2).
 - Governance. India's AI Governance Guidelines (November 2025) favour a light-touch, sector-led approach without an umbrella AI law (Ministry of Electronics and Information Technology, 2025). That approach can be adapted quickly, but its effectiveness depends on sectoral regulators' capacity.
 - Strategy revision. The foundation-model shift was absorbed through a mission rather than a revised national strategy with a built-in review cycle.
-

8. Recommendations: Shortening the Next Lag

The first version proposed a 90-day emergency reset. Since the funded response has now happened, the more durable question is how to make the next lag shorter. Five institutional mechanisms follow from the hypotheses in Section 5:

1. Scheduled strategy review. Revise the national AI strategy on a fixed cycle (for example, every two years), with a published record of what changed and why (H1, H5).
 2. Paradigm-shift triggers. Define in advance the indicators that force an out-of-cycle review, such as a large capability jump on public benchmarks or a sharp fall in training costs (H1).
 3. Flexible funding windows. Reserve a share of mission funds for rapid reallocation within a year of a trigger, without requiring a new approval cycle (H4).
 4. Compute for academia as standing infrastructure. Maintain national compute access for universities as permanent infrastructure, so that academic research can move with the field (H2).
 5. Return pathways for talent. Pair compute and funding with positions designed to attract researchers in the diaspora back to Indian labs, and measure net migration annually (Section 4.2).
-

9. Limitations

- Choice of signals and responses. Lag values depend on which signal and which response are chosen; the paper states its choices and reports alternatives, but others are defensible.
 - Different kinds of response. Regulatory, industrial and strategic responses are not directly comparable (Section 3.4).
 - Associational evidence. The indicators in Section 4 coincide with the lag; the paper does not establish that the lag caused them.
 - Source quality. Some indicators come from industry analyses, government press releases and news reports rather than peer-reviewed studies; these are identified as such.
 - Moving target. India's AI policy and model landscape changed rapidly in 2025–2026, and parts of this analysis may date quickly.
-

10. Conclusion

India did not stand still between 2018 and 2024, but it did not respond at the national level to the most important change in AI of the period until about four years after specialists could see it. When it did respond, it moved fast: compute beyond target within two years, sovereign models within three. That speed is itself the strongest evidence that the earlier delay was a matter of decision, not capacity. The costs of the lag are visible in research share and above all in the net outflow of AI talent, and they will not disappear because the lag has ended. The lesson is institutional. Paradigm shifts in AI will keep coming,

and the countries that respond fastest will be those that decide in advance how they will notice and how they will change course.

References

BharatGen. (2025). *PARAM-1 BharatGen 2.9B model* (arXiv:2507.13390) [Preprint].
<https://arxiv.org/abs/2507.13390>

Brown, T., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

Business Standard. (2026, February 19). Why Sarvam's new 105B model marks a shift in India's sovereign AI ambitions. https://www.business-standard.com/technology/tech-news/sarvam-105b-model-sovereign-ai-india-foundation-model-launch-impact-summit-126021900551_1.html

Change Engine. (2024). *India's share of papers at top AI conferences, 2018–2023* [Industry analysis]. Reported in: India ranks far below US, China in AI research with only 1.4% paper contribution to conferences. *The Wire*. <https://m.thewire.in/article/tech/india-ranks-far-below-us-china-in-ai-research-with-only-1-4-paper-contribution-to-conferences>

Cyberspace Administration of China, et al. (2023). *Interim measures for the management of generative artificial intelligence services* (issued 10 July 2023, effective 15 August 2023). English translation: China Law Translate. <https://www.chinalawtranslate.com/en/generative-ai-interim/>

DeepSeek-AI. (2024). *DeepSeek-V3 technical report* (arXiv:2412.19437) [Preprint].
<https://arxiv.org/abs/2412.19437>

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.

Executive Office of the President. (2023). Executive Order 14110: Safe, secure, and trustworthy development and use of artificial intelligence. *Federal Register*, 88(210), 75191–75226.

Gala, J., Chitale, P. A., et al. (2023). IndicTrans2: Towards high-quality and accessible machine translation models for all 22 scheduled Indian languages. *Transactions on Machine Learning Research*.
<https://arxiv.org/abs/2305.16307>

Ministry of Electronics and Information Technology. (2025). *India AI governance guidelines: Enabling safe and trusted AI innovation*. Government of India.
<https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf>

NITI Aayog. (2018). *National strategy for artificial intelligence #AIforAll*. Government of India.

NITI Aayog. (2020). *AIRAWAT: Establishing an AI specific cloud computing infrastructure for India*. Government of India. <https://www.niti.gov.in/sites/default/files/2023-03/AIRAWAT-Establishing-an-AI-Specific-Cloud-Computing-Infrastructure-for-India.pdf>

NITI Aayog. (2021). *Responsible AI #AIforAll: Approach document for India, Part 1: Principles for responsible AI*. Government of India. <https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf>

Press Information Bureau. (2024, March 7). *Cabinet approves ambitious IndiaAI Mission to strengthen the AI innovation ecosystem*. Government of India. <https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012355>

Press Information Bureau. (2025, December 30). *Transforming India with AI: Over ₹10,300 crore investment and 38,000 GPUs powering inclusive innovation*. Government of India.

State Council of the People's Republic of China. (2017). *New generation artificial intelligence development plan*. English translation: DigiChina, Stanford University. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>

Stanford Institute for Human-Centered Artificial Intelligence. (2026). *The AI Index 2026 annual report*. Stanford University.

The White House. (2025). *America's AI action plan*. Executive Office of the President.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.