CHAPTER 14

EVOLUTION OF THE AI ECOSYSTEM IN INDIA: THE WAY FORWARD

This chapter examines how Artificial Intelligence (AI) is reshaping the global economy and outlines a pragmatic strategy for India in an environment marked by rapid technological change and persistent uncertainty. The chapter argues that India's AI strategy should be grounded in its own economic realities, rather than attempting to replicate unsustainable models adopted in advanced economies.

Given the constraints related to capital, computing capacity, energy, and infrastructure, pursuing scale for its own sake is neither efficient nor necessary. Instead, the chapter makes the case that a bottom-up, multiple sector-specific approaches under a single vision has the potential to pay dividends and turn into a source of dignified employment for India's youth. India's development of AI must be grounded in open and interoperable systems to promote collaboration and shared innovation. This pathway aligns more closely with India's strengths in human capital, data diversity and institutional coordination.

On governance, the chapter emphasises sequencing, enabling experimentation first, scaling next, and introducing binding obligations only where risks and asymmetries are most pronounced. The proposed framework for data governance strikes a balance between openness to cross-border flows and strengthening accountability and regulatory visibility. It is rooted in the objective of ensuring that the value accruing from India's domestic data is retained within the country for the benefit of the people. The government's role is framed as that of an enabler and coordinator, helping markets and institutions adjust in step with technological change.

Overall, the chapter treats AI as a strategic choice. The central message is that India's opportunity lies in deploying AI in a way that is economically grounded and socially responsive.

ARTIFICIAL INTELLIGENCE IN INDIA'S ECONOMIC CONTEXT

Introduction: What we know, and what we still don't

14.1. When the Economic Survey last examined Artificial Intelligence (AI) in early 2025, the conversation, both globally and in India, was still dominated by the possibilities and potential of AI. While AI was already visible in some areas as a productivity-enhancing tool or embedded within some service platforms, its broader economic implications remained largely conjectural. The focus of the chapter, therefore, was appropriately placed on preparedness, discussing skills, data, infrastructure, and institutional readiness.

14.2. One year on, the conversation has definitely shifted. AI is no longer a distant or speculative technology. It is increasingly being adopted, even if in an experimental capacity, in organisations around the world. Based on a survey of 1993 firms by McKinsey, 88% of organisations surveyed in 2025 reported that they are utilising AI in at least one of their business functions.[1] Of those using it, 31% are in the process of scaling its application across the organisation, while 7% have already fully deployed and integrated AI.

Chart XIV.1: Use of AI for at least one business function by respondents' organisations, McKinsey

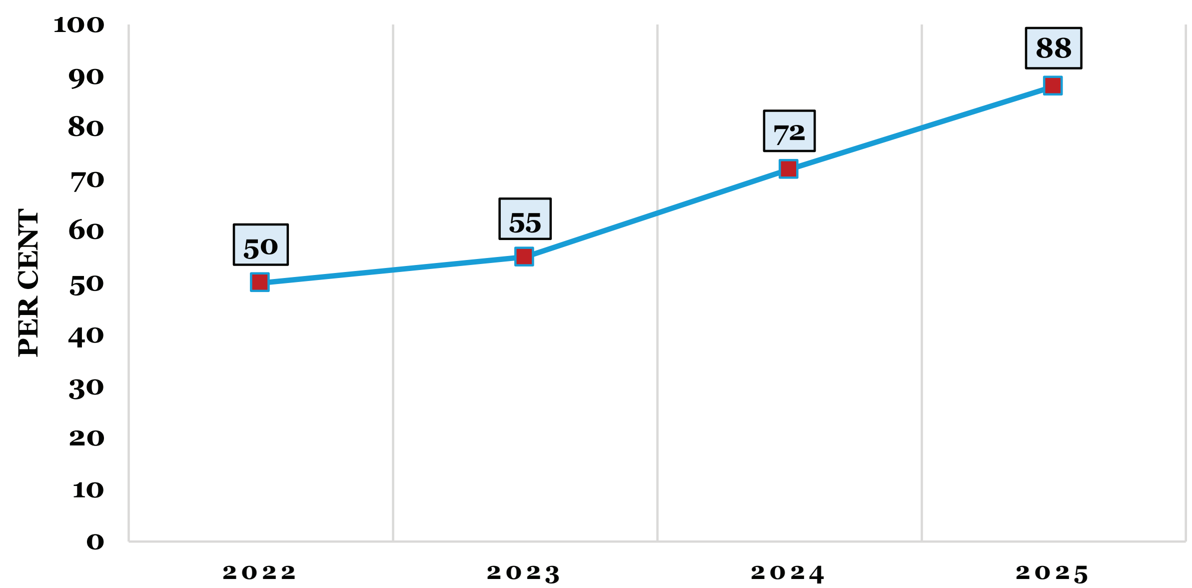

Source: Exhibit 1, Page 4, The State of AI in 2025 Report, McKinsey. https://tinyurl.com/msmk4x3s.

14.3. Although utilisation is currently concentrated in High-Income countries (accounting for 58.4% of all usage in April 2025), utilisation in Upper- and Lower-Middle-Income countries has also expanded (22.5% and 18.7%, respectively).[2] Innovations and continuous improvement in AI capabilities are driving firms and new start-ups to develop ways in which AI can be applied to solve real-world problems.

1 The State of AI in 2025 Report. McKinsey. https://tinyurl.com/msmk4x3s.

2 Digital Progress and Trends Report: Strengthening AI Foundations. World Bank. https://tinyurl.com/3k9hebmx.

14.4. At the same time, greater visibility into AI adoption has also brought greater clarity on the nature of the technology itself. Over the past year, it has become evident that while the use of AI tools can be widespread, the frontier of AI remains highly concentrated. The development and training of advanced foundational models is increasingly capital-, compute-, data- and energy-intensive, favouring a small set of firms with access and the political capital to secure large-scale infrastructure projects, specialised hardware, and deep pools of technical talent.

14.5. Early evidence has also begun to temper some of the more extreme predictions surrounding AI's near-term labour impact. For instance, a study conducted by Yale's Budget Lab indicates that the broader labour market in the United States has not experienced a discernible disruption due to AI.[3] Similarly, a study by Brynjolfsson, Chandar, and Chen (2025)[4] highlights that the difference in job prospects between occupations highly exposed to AI and those with relatively low exposure is minor. According to Renault (2025), most Danish workers also benefit from the adoption of AI[5]. The emerging evidence does provide some reassurance in the near term, especially for labour-abundant economies such as India.

14.6. However, this does not invite complacency, especially from a policymaker's perspective. While labour may be complemented in the near term as organisations work to incorporate AI into their tasks, productivity gains from augmentation have a ceiling, as highlighted in Box XIII.2 of the Economic Survey 2024-25. Further, Box XIV.1 demonstrates a meaningful structural shift in the interaction between employment and output, using the United States as a case study. All in all, caution is still warranted as India attempts to solve the puzzle of AI and labour. This represents one of the most considerable looming uncertainties about the technology.

Box XIV.1: Employment in the AI-Salient Period - Evidence from the U.S. Service Sector

Public debate on the labour market implications of AI is highly polarised. One camp predicts large scale displacement while the other side notes that the effect of AI on labour markets will remain muted for the foreseeable future. These views are often framed in terms of labour substitution vs augmentation, yet empirical evidence capable of distinguishing between these outcomes remains scarce. The reason no definitive conclusions have been reached on anything AI-related is twofold. One, AI is relatively new and integrating it into existing

3 Gimbel, M., Kinder, M., Kendall, J., & Lee, M. (2025). Evaluating the Impact of AI on the Labor Market: Current State of Affairs. Technical report, The Budget Lab at Yale University. URL: https://budgetlab. yale. edu/research/evaluating-impact-ai-labor-market-current-state-affairs.

4 Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Digital Economy.

5 Renault, T. (2025). The Impact of Artificial Intelligence on Denmark's Labor Market. https://doi.org/10.5089/9798229021647.018.

economic structures is not straightforward. Two, considering that AI is still undergoing constant development, there is considerable uncertainty surrounding its capabilities and costs. Enterprises want to take risks and adopt AI, but also avoid being the ones who get it wrong.

However, the core uncertainty is not whether AI matters for this sector, but how and when it will manifest. Technological transitions do not arrive overnight as instantaneous shocks. Evidence from past automation waves suggests that early-stage adoption complements labour, with substitution pressures only emerging once the market saturates and productivity gains no longer lead to significant cost reductions.

Against this backdrop, our analysis examines employment dynamics in the United States' professional, business and information service sector, focusing on non-supervisory private-sector workers. The United States is used as a case study because its white-collar and services-heavy labour market sits closest to the technological frontier of AI adoption. Any dynamics observed can offer early-signals and policy-relevant insights for India as its own services-led economy moves along a similar, albeit lagged, trajectory.

Using a time-series framework, the analysis evaluates whether the period coinciding with the widespread introduction of Generative AI tools is associated with a change in the relationship between employment, output, wages, and financial conditions. Our analysis is explicitly framed as an assessment of structural change rather than a test of causal effects.

Estimation and Findings

An Autoregressive Distributed Lag model was estimated with nonsupervisory employment in the Professional, Business, and Information Services (PBIS) sector as the dependent variable for the period from March 2016 to July 2025. Explanatory variables include the average hourly wages in the sector, the federal funds rate to account for the cost of borrowing, a monthly real GDP series created by S&P, and the consumer price index. To assess potential structural change in employment dynamics, the model incorporates a post-December 2022 regime binary dummy, corresponding with the period of heightened salience of GenAI tools, along with an interaction between this regime indicator and real output.

Our estimates indicate a change in the employment dynamics of the PBIS sector during the post-December 2022 period. Rather than indicating a discrete break in employment levels, the estimates point to a change in how employment responds to output growth in the post-2022 period.

Conditional on wages, output, and the interest rate, the post-December 2022 regime indicator is positively associated with PBIS employment, both in the short- and long-run. This association is interpreted as a structural shift in employment dynamics during the AI-salient period, rather than as evidence of direct causal effect of AI adoption. At the same time, the interaction between the post-2022 regime indicator and real GDP is negative and statistically significant across both time-horizons. This implies that relative to pre-2022, the marginal responsiveness of PBIS sector employment to output growth is lower in post-2022 regime.

Sectoral Contrast and Falsification Exercise

To assess whether the post-2022 regime indicator captures a sector specific structural shift rather than a broader macroeconomic phenomenon, an alternative specification was estimated using private sector non-supervisory employment that excludes the PBIS sector. The results reveal a stark difference.

In the non-PBIS sector, the post-2022 indicator is associated with a modest short-run coefficient and a near zero, statistically insignificant long-run coefficient. This stands in contrast to the PBIS estimates, where the corresponding short-run coefficient is substantially larger and the long-run association remains statistically significant.

The divergence suggests that the structural change identified in the PBIS sector is not a general feature of post-pandemic labour market dynamics, monetary conditions, or aggregate demand recovery. Instead, it appears to be concentrated in sectors characterised by higher sensitivity and exposure to GenAI tools.

To reiterate, this does not imply that GenAI has directly affected employment in PBIS while leaving others untouched. Rather, it demonstrates that the PBIS sector is exhibiting a markedly distinct adjustment pattern. The sectoral contrast strengthens the interpretation of the post-2022 dummy as capturing a meaningful structural shift rather than a purely cyclical or economy-wide effect.

Discussion

The results suggest that employment adjustment in the PBIS sector is more nuanced than the polarised debate around AI-driven job losses would imply. Rather than indicating an abrupt contraction in employment, the estimates point to a change in how employment responds to output growth in the post-2022 period.

In this context, the positive association between the post-2022 indicator and PBIS employment is best interpreted as reflecting organisational and structural adjustments during a period of heightened technological salience. Firms may be expanding or reallocating labour to integrate new systems, manage workflows, and redesign service delivery, consistent with earlier findings that productivity-enhancing technologies can initially complement labour in sectors with elastic demand (Bessen, 2019[6]; Albanesi et al., 2024[7]). This may also be why a recent study published by Yale[8] found no net-negative effects of AI on labour demand in the United States.

Then why do many experts focus so much on automation and the potential for AI-induced displacements? Why are policymakers concerned about the adverse consequences of AI diffusion? This can be explained by the coefficient of the interaction term included in our model. The negative interaction between the post-2022 indicator and output suggests that marginal responsiveness of employment to economic growth has weakened relative to the pre-2022 period. Or put simply, the labour intensity of output has marginally declined. While this does not constitute evidence of task-level automation, it does indicate a shift in aggregate employment elasticities under a new technological regime.

6 Bessen, J. (2019). Automation and jobs: When technology boosts employment. Economic Policy, 34(100), 589-626, https://tinyurl.com/2c7zmj3c

7 Albanesi, S., Dias da Silva, A., Jimeno, J. F., Lamo, A., & Wabitsch, A. (2024). New technologies and jobs in Europe. Economic Policy, eiae058, https://tinyurl.com/ymarmumr

8 https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Taken together, the estimates imply that policymakers and firms might witness a non-linear labour market trajectory. Unless the labour market adapts and new skills are learned, which alter the profile and types of jobs people are engaged in, we may observe even more reductions in the labour-intensity of GDP in the future. Paraphrasing the words of T.S. Elliot, the change comes not in a single shock, but in a quiet, steady drift.

Table XIV.1: Short-run dynamics[9]

Variable	PBIS Sector (Std. Error)	Non-PBIS Sector (Std. Error)
D(Regime Indicator)	25.74 (3.74)	1.58 (0.11)
D(Interaction)	-2.572 (0.37)	-0.00007 (-0.00179)
Error Correction Term	-0.217 (0.04)	-0.07 (0.009)
Source: Author's Calculations		

Table XIV.2: Long-Run Dynamics

Variable	PBIS Coeff (Std. Error)	Non-PBIS Coeff.
AI	9.409 (1.556)	1.90 (1.92)
Interaction	-0.938 (0.154)	-0.00795 (-0.00819
Source: Author's Calculations Bounds Test F-Statistic is significant at 5%		

14.7. Uncertainty also extends into the evolving structures of the global AI ecosystem. Control over critical inputs, such as data, compute, models, and standards, is increasingly concentrated. This raises concerns about market power, technological dependence, and the resilience of supply chains. It also raises a substantial question about the future of India's IT sector, as firms that once relied on India's comparative advantage to handle a bulk of their work may no longer need to do so. It risks hollowing out India's core value proposition if adaptation lags. If the country is to sustain its competitive edge in IT, a comprehensive evolution is necessary, one that takes full advantage of the potential embedded in AI development and deployment.

9 Coefficients of other macro variables are not included here for brevity.

14.8. Another source of uncertainty pertains to defining a regulatory approach to AI. Countries diverge in how they design their institutions to address AI-related challenges, reflecting differing priorities. For advanced economies, their decisions will determine how AI can be leveraged to enhance productivity in the face of labour shortages and an ageing population. For India, the challenge is to govern AI in a manner that is sensitive to its economic realities. The choices that India's own institutions make will play a central role in determining not only the pace of AI diffusion but also how its economic value is distributed across sectors and among people.

14.9. The question, therefore, is how India positions itself within an ecosystem marked by rapid technological change, concentrated capabilities, limited resources and persistent uncertainty. Will it remain primarily a consumer of AI technologies or emerge as a meaningful contributor to their development and deployment? This is what the chapter aims to contribute towards.

14.10. In the sections that follow, we outline the trade-offs that India will have to contend with and how the nation can evolve its AI ecosystem in a manner that best aligns with its goals. The following sub-section details the trade-offs and asymmetries involved in building an AI ecosystem. Sections 2 to 4 highlight the strategic necessity of India's own AI solution and propose practical steps that can be taken to evolve the ecosystem, including AI model and application development, human capital, and how governance needs to evolve, including a framework for incentivising data localisation. Section 5 highlights the various safety concerns that policymakers must be aware of. Section 6 puts forth a phased deployment proposal, and concludes.

Asymmetries and Trade-offs in the AI Ecosystem

14.11. The AI ecosystem is characterised by pronounced asymmetries across countries, firms, capabilities, and stages of the value chain. These asymmetries are not incidental; they are structural outcomes of how AI is financed, developed and deployed. For India, recognising these asymmetries is essential, as policy choices in AI are constrained by the trade-offs these asymmetries present to the country.

Frontier versus Application: Capability Asymmetries

14.12. At the core of the global AI divide lies a sharp distinction between frontier model development and application-led development. While AI usage has been growing, as highlighted by the World Bank report cited earlier, the capability to design and train large foundational models remains highly concentrated in the hands of a few large firms. These firms exercise significant control over the market and exert high demand pressures on the resources necessary for AI, allowing them to erect high barriers to entry.

14.13. Add to this the export restrictions imposed on the most advanced processors required for scaling up frontier model development, and the task ahead of India becomes extremely challenging. This creates a fundamental asymmetry: most countries may end up participating in AI primarily as users, while a few will shape the technology's trajectory, standards, cultural leanings, and pricing. Attempting to close this gap would involve prohibitive fiscal costs towards what is increasingly becoming an unsustainable approach to AI development. The trade-off, therefore, is between expending scarce resources to chase frontier-scale models or deploying those resources more effectively towards domain-specific AI systems aligned with domestic economic priorities.

Scale versus Inclusion: Capital-Labour Trade-offs

14.14. AI adoption alters incentives within firms by raising the marginal productivity of capital relative to labour, particularly in white-collar service sectors. Firms driven by cost reductions and maximising productivity gains are more inclined to rapidly scale up AI adoption[10], increasing the likelihood of capital-labour substitution in specific task categories. Most of these tasks are concentrated in the low-value-added segments of the service sector, and high exposure to AI[11] means firm leadership may view these jobs as redundant.

14.15. For labour-abundant economies such as India, this creates a tension between aggregate productivity gains and employment absorption. Rapid, uncalibrated deployment of AI may boost output but risks displacing segments of the workforce faster than the economy can reabsorb them. Conversely, delaying adoption to protect jobs may risk locking firms into a low productivity equilibrium. The policy challenge, therefore, is not whether to adopt AI, but how to pace its diffusion so that labour augmentation can be facilitated.

Open versus Proprietary Models: Cost, Control, Dependence

14.16. Another asymmetry lies in the ownership and governance of AI models. The most widely deployed AI models are proprietary and opaque, limiting transparency around their training data, internal logic, and update mechanisms. Since these models are essentially black boxes, users can never know what changes are being made to the underlying source code and how it might be altering the behaviour of the AI model.

14.17. In contrast, open-source models and open-weight models offer lower entry barriers, greater adaptability, and reduced vendor lock-in. However, they also include trade-offs

10 From Financial Times. Accessed 23rd December 2025. https://tinyurl.com/mr42snkj.

11 Tomlinson, K., Jaffe, S., Wang, W., Counts, S., & Suri, S. (2025). Working with AI: measuring the applicability of generative AI to occupations. arXiv preprint arXiv:2507.07935.

related to quality control and fragmentation if not guided by a coherent national strategy. India's challenge is to strike a balance between openness and stewardship, leveraging shared innovation while ensuring that the economic value created from domestic data and intellectual property accrues within India rather than being captured abroad.

Compute Intensity versus Resource Constraints

14.18. AI development is inextricably linked to physical infrastructure. Data centres require large quantities of electricity[12] and water[13], and AI workloads introduce volatility into power demand, posing risks to grid stability[14]. Global experiences show that AI-driven data centre expansion can place significant strains on existing energy systems, even in advanced economies.

14.19. In a recent Congressional testimony, Mark P. Mills, senior fellow at the Manhattan Institute noted, "*The unprecedented digital construction underway has led to the rediscovery of a basic truth: All software exists inside hardware that, in turn, uses energy, a lot of it. Each digital byte uses an infinitesimal amount of energy. But here we find salience for the euphemism that quantity has a quality all its own. Again, in moonshot terms: the amount of energy used to launch a rocket is consumed every day by just one AI-infused datacenter.*"[15]

14.20. Similarly, the CEO of IBM raised concerns about the financial viability of the capital expenditures being undertaken for data centre expansion[16]. With some firms projected to burn half a trillion in cash by 2030 while pursuing compute infrastructure

12 To grasp the scale of electricity demand from AI data centres, one need only look at the recent revision by the U.S. Energy Information Administration (EIA). Previously, the EIA had projected a modest 2 % annual growth in electricity consumption through 2026. In light of the rapid expansion of AI-driven data centre activity, the forecast has been sharply revised upwards, to 3 % growth in 2025 and 5 % in 2026. See Short-Term Energy Outlook, June 2025. US Energy Information Administration.

13 AI Data Centers consume upto 20 lakh litres of water per day, and globally, they consume 56,000 crore litres of water annually (as recently reported by Bloomberg, https://tinyurl.com/wsdmmbjr). If India scales up AI Data Centers, it has the potential to add an extraordinary amount of stress on our already strained groundwater and freshwater reserves. Also see analysis from the University of Illinois (https://tinyurl.com/yc5fax8j), Stanford University (https://tinyurl.com/bbce9en5), and the New York Times (https://tinyurl.com/mr283x7f).

14 For instance, an incident on 10th July 2024 in Northern Virginia demonstrated how the simultaneous loss of 1500 MW of load from the data centres caused significant voltage depression and frequency deviation in the power grid. Swift intervention by the grid operators prevented a massive surge in the electrical systems, only narrowly avoiding a total blackout in the region. The incident demonstrates how the proliferation of AI Data centres across India can add significant strain to the grid. The integration and power usage of AI Data Centres is a topic that requires separate and serious consideration. See North American Electric Reliability Corporation Incident Review. Published on 8th January 2025. https://tinyurl.com/2nrdfa48.

15 Testimony of Mark P. Mills, Executive Director of the National Center for Energy Analytics Before Subcommittee on Economic Growth, Energy Policy, and Regulatory Affairs U.S. House Committee on Oversight and Government Reform. https://tinyurl.com/y5zjymaa.

16 IBM CEO questions profitability of AI Data Centre Boom. https://tinyurl.com/bdfuwfdc.

creation for AI[17], the risks of a financial contagion[18] stemming from the unprecedented scale of a debt fuelled expansion[19] remains high.

14.21. For India, as power, finance, and especially water remain binding constraints, scaling compute indiscriminately carries opportunity costs. Investment in AI infrastructure competes directly with other sources of demand, such as households and industries. This creates a trade-off between centralised scale and distributed efficiency, strengthening the case for smaller, task-specific models that can run on limited hardware and decentralised compute networks.

Regulation versus Innovation

14.22. An age-old trade-off that policymakers must always contend with is the balance between sustaining innovation and designing regulatory frameworks intended to keep the broader public safe. Furthermore, the trade-off between the two is inherently asymmetric across countries. Regulatory compliance, aimed at ensuring safety, controlled proliferation, auditing, transparency, or mitigating liability exposure, imposes fixed costs that may scale poorly for smaller firms and those involved in early-stage experimentation.

14.23. Frontier AI ecosystems in high-income countries and other cash-rich large firms often demonstrate a capacity to absorb these costs. India's more fragmented and resource-constrained innovation landscape could be stifled if the same degree of regulations were binding.

14.24. However, very minimal or completely absent regulatory clarity can undermine trust, slow adoption, and create systemic risks, particularly as AI is deployed in critical sectors such as healthcare, education, governance, or finance. The challenge for India, therefore, is how and when to regulate AI.

Strategic Autonomy versus Global Integration

14.25. Lastly, AI has emerged as a geostrategic asset. Export controls on advanced chips, restrictions on technology transfer, and the weaponisation of software and inputs necessary to build AI infrastructure have underscored the risks of overdependence on foreign systems. As AI applications diffuse into critical sectors and public institutions, these dependencies carry systemic risks.

17 OpenAI needs to raise atleast USD 270 bn by 2030 so it can continue to lose money. Financial Times. Accessed on 27 December 2025. https://tinyurl.com/mpkk86p5.

18 IMF World Economic Outlook, October 2025. https://tinyurl.com/ywkrzt73.

19 Tech groups shift USD 120 bn of AI data centre debt off balance sheets. Financial Times. Accessed on 27 December 2025. https://tinyurl.com/3ywmvraf.

Chart XIV.2: Over 70% of all data centres (by count) are located in High-Income Countries (as of June 2025)

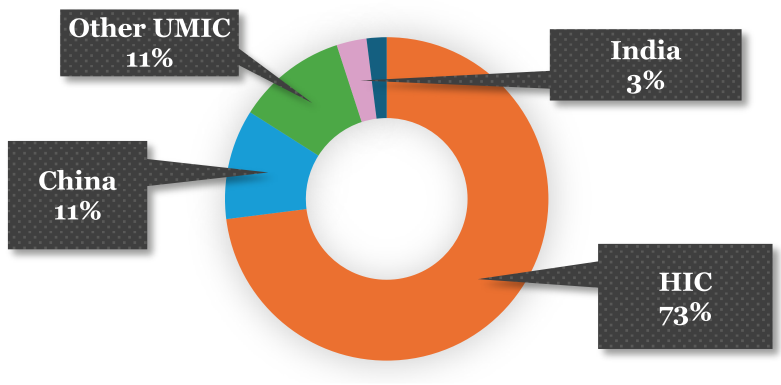

Source: Digital Progress and Trends Report: Strengthening AI Foundations. World Bank. https://tinyurl.com/yhr9hav7.

14.26. Yet, complete technological self-sufficiency is neither feasible nor efficient. The trade-off, therefore, lies between strategic autonomy and continued integration with global innovation networks. India's policy choices must navigate this balance carefully, preserving openness where it enhances capability while insulating critical functions from external shocks.

14.27. These asymmetries do not imply that India is at a structural disadvantage in the AI era. Instead, they define the constraints within which a viable and sustainable AI strategy must be formulated. The sections that follow build on this diagnosis to outline pathways that seek to maximise economic and social returns from AI while remaining mindful of the trade-offs.

A DEVELOPMENT-ORIENTED APPROACH TO AI

The necessity of India's own AI solution

14.28. The AI ecosystem remains sufficiently young that existing imbalances and constraints need not define its future trajectory. This creates an opening for India to shape a more value-creating and dignified employment opportunity for its workforce. As NITI Aayog's report highlights, there are numerous potential benefits to indigenous AI development[20]. Multinational IT companies have thus far benefited from "low-cost labour" in India, but AI presents an opportunity to shift perceptions and structurally change the reality of India's labour market.

14.29. Secondly, developments in AI will have natural spillovers into existing sectors such as services, manufacturing, defence and power. In this context, reliance on foreign

20 AI for Viksit Bharat: The Opportunity for Accelerated Economic Growth. NITI Aayog.

multinationals for AI-based solutions will leave India vulnerable to shifts in geopolitics, potentially constraining the country's future diplomatic choices. Just as critical minerals and semiconductors are utilised to shape foreign policy, AI capability and resources will similarly be utilised for geostrategic negotiations.[21]

14.30. Accordingly, AI should not be regarded merely as a technological advancement, but as a strategic priority with far-reaching implications for the future of India's critical infrastructure, labour market, foreign policy and culture.

Box XIV.2: Identifying Bottlenecks to AI Compute Expansion - An Agent Based Model Approach

As discussed earlier in the chapter, the availability of sufficient compute capacity in India is a necessary condition for the training and development of cutting-edge models. However, surging demand for inputs (driven by higher GPU demand) and constrained supply conditions (due to shortages in the availability of high-bandwidth memory chips and storage) are driving cost concerns[22]. This, in turn, is bound to have a ripple effect on the cost of expanding compute capacity in India making financial viability a possible bottleneck.

Similarly, higher demand from buyers abroad can choke the available supply of GPUs[23]. Thus, even if the sovereign, domestic investors, or financial institutions are willing to finance data centre expansion, plans may need to be put on hold until GPU supplies are secured. In this context, the objective of this exercise is to examine how the expansion of data centre and AI compute capacity is shaped by the interaction of financial constraints, infrastructure readiness, and external supply dependencies.

To do so, the exercise adopts an Agent Based Modelling (ABM) approach where financing, grid availability and hardware access are allowed to interact and shape one another through feedback loops. The focus is on identifying conditions under which coordination failures emerge, even in the presence of strong underlying demand.

The exercise is intended as a policy stress test rather than a forecasting tool. It does not aim to estimate optimal capacity levels or precise investment requirements. Instead, it is designed to surface structural vulnerabilities in the AI infrastructure ecosystem, identify leverage points where policy interventions may be most effective, and clarify the limits of price-based or subsidy-led approaches when financial and infrastructural constraints are binding, as is the case for India.

21 Dr. Chris Miller's book titled "Chip Wars" offers a sobering reality check about our dependence on a handful of nations for semiconductors which are extremely critical to every facet of the modern economy. The control over certain processes within the supply chains affords nations a very powerful bargaining tool for shaping global geopolitics in their own favour. AI is expected to be used in a similar way.

22 NVIDIA in its Q3 2026 Earnings Call expressed concern over the rising cost of inputs. https://tinyurl.com/mtjc88tz.

23 OpenAI recently signed a deal with two of the only three major high-bandwidth memory manufacturers in the world to purchase 40% of the global supply of memory chips. https://tinyurl.com/368u9kkd.

The simulation was run for a period of 40-quarters, capturing the evolution of constraints over the medium-term.

Scenario 1: Baseline

For the baseline scenario, demand for data centre capacity is assumed to grow at a rate of 24% per annum, which is on the conservative side of expectations[24]. Power tariffs start at ₹8 per kWh, and are elastic within the system. Additional power capacity is added to the grid with a lag of 12 quarters and the cost of borrowing for data centre operators at baseline is set at 9% per annum. Based on the above configuration, the system is setup and simulated. Letting the model run for 40 quarters, in discrete time, reveals the outcomes of our baseline scenario.

At base, the model exhibits a sequencing of bottlenecks over time, rather than a simultaneous binding of constraints. In the early quarters, the dominant friction is access to finance, as operators attempt to scale capacity ahead of realised revenues. Despite strong and steadily growing domestic demand, balance-sheet constraints and investor hurdle rates initially prevent a subset of firms from expanding, leading to a short-lived spike in entities waiting for finance. As revenues accumulate from the rising cost of compute and margins stabilise, this financial bottleneck eases, indicating that under normal macro-financial conditions, financing constraints are not permanently binding but act as an initial filter on expansion timing.

Chart XIV.3: Scenario 1 - Baseline

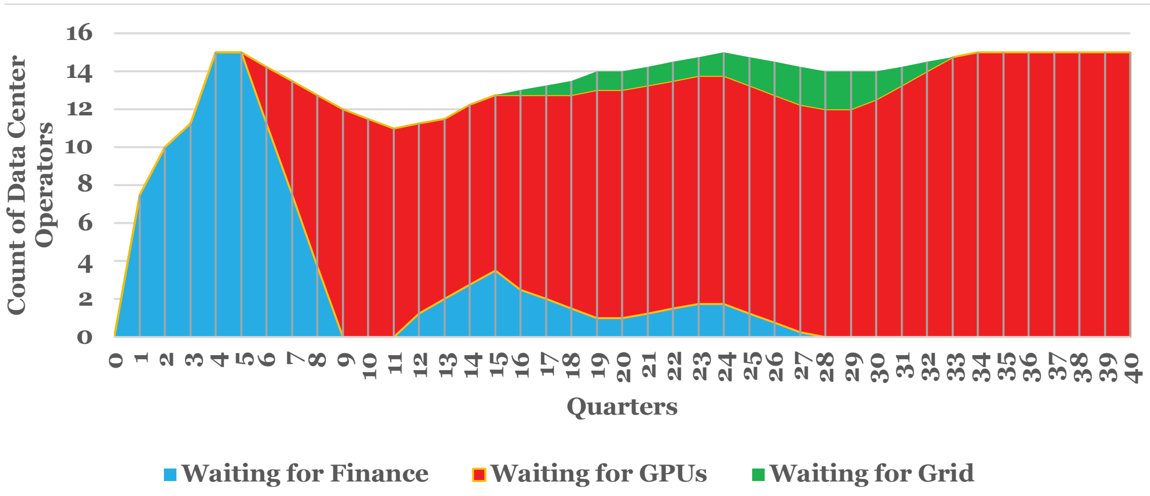

Source: Author's Calculations

As the simulation progresses, the primary constraint shifts toward GPU availability. Even with "normal" global supply conditions and India's GPU share in global demand set at a realistic 4%, the probabilistic access mechanism and minimum lead times result in a persistent and growing pool of operators waiting for hardware. This bottleneck dominates for most of the simulation horizon, suggesting that hardware access, rather than pricing or power availability, becomes the central limiting factor in sustained capacity expansion.

24 See, for instance, projections by S&P. https://tinyurl.com/2fwhrkmu.

Grid-related constraints emerge only episodically but remain secondary. Overall, the baseline outcome illustrates a system where demand is not the binding constraint; instead, expansion is paced by the slowest-moving complementary inputs, with hardware supply uncertainty exerting the longest lasting drag on scale-up.

Scenario 2: Elevated Foreign GPU Demand (Relative to Baseline)

In the second scenario, foreign demand is set to expand at almost double the pace of the baseline scenario, with all other factors remaining the same. Running the simulation for 40 quarters once again sheds light on some interesting dynamics. Higher foreign demand directly impacts the price of GPUs, worsening project economics by increasing upfront capex required and challenging the financial viability of data centre projects in India. Higher prices also weigh down on the profitability margins and the debt service capacity of data centres, making investors and banks reluctant to finance further expansion. The financial bottleneck remains a binding constraint for significantly longer than in the baseline.

Chart XIV.4: Scenario 2 – High Foreign Demand for GPUs

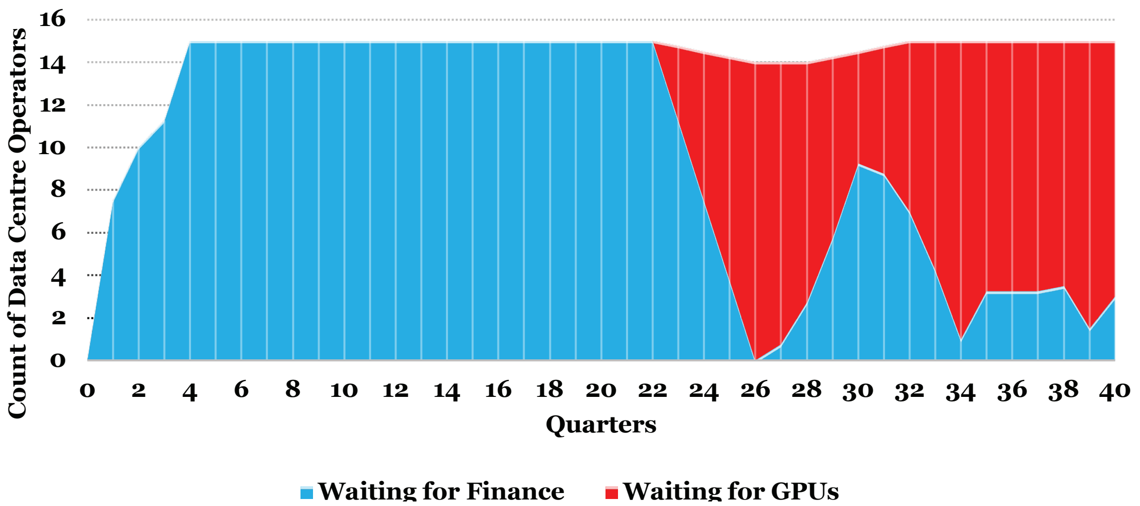

Source: Author's Calculations

Over time, even as a subset of operators gradually overcome initial financing constraints due to the rising cost of compute contributing to increased revenues, uncertainty around global GPU supply continues to bind. While balance sheets improve and access to finance partially normalises, elevated global demand keeps GPU availability probabilistic and delivery timelines uncertain. This creates a second-order constraint: projects that are financially viable on paper are still unable to proceed due to unresolved hardware access. The result is a decoupling between financial readiness and execution capability, underscoring that easing domestic financial constraints alone is insufficient when hardware supply remains externally constrained.

Scenario 3: High Domestic Demand with liberal financing (Relative to Baseline)

In scenario 3, foreign demand is returned to baseline but India's demand is now set to grow at 32% per annum. India's share in global GPU demand is increased to 6% a year. To ensure there is sufficient finances available to fund the growing appetite, investor hurdle rates are lowered to 10% (relative to 15% for the baseline scenario), and the financial system is flush with liquidity. This reduces the incidence and duration of financing bottlenecks in scenario 3, with a lesser number of operators waiting for a decreased period of time (2 less quarters relative to baseline) to avail funding.

Chart XIV.5: Scenario 3 – High Domestic Demand with liberal finance availability

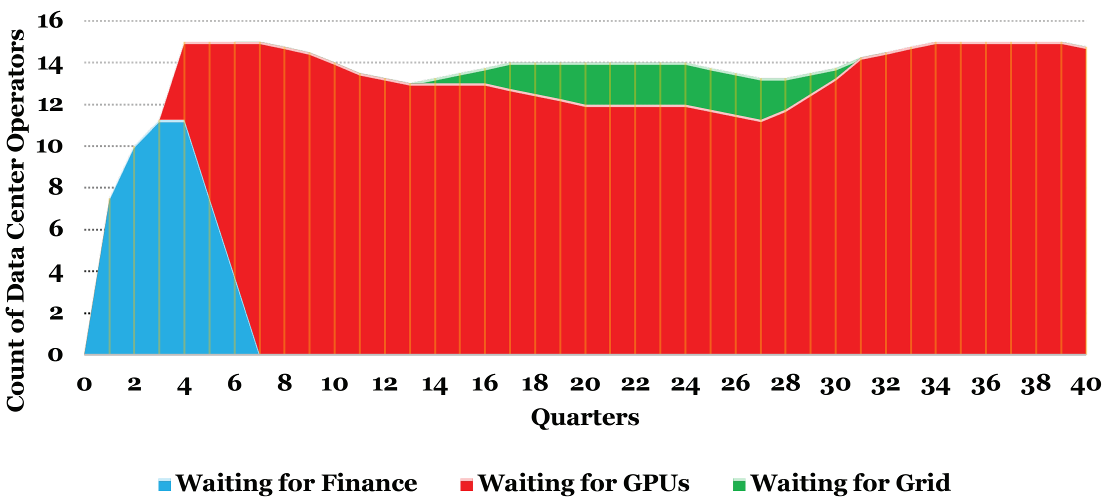

Source: Author's Calculations

However, the relaxation of financial constraints does not translate into proportionate capacity addition. Instead, the dominant bottleneck shifts decisively to GPU availability, with grid constraints appearing intermittently as utilisation rises. Higher domestic demand accelerates the pace at which projects reach the execution stage, but hardware access remains governed by global supply conditions. As a result, a large share of operators accumulate in the "waiting for GPUs" state for much of the simulation horizon. Compared to the baseline, the system moves faster into this hardware-constrained regime, while compared to Scenario 2, the constraint emerges despite favourable global conditions, underscoring that abundant capital and strong demand are insufficient to overcome physical and geopolitical limits in advanced hardware supply chains.

Concluding Thoughts

Across all scenarios, a recurring pattern is that developments in global GPU supply chains play a meaningful role in shaping the long-run pace of AI infrastructure expansion. This persists even when domestic demand conditions are strong and financial constraints are eased, highlighting an important consideration for India's AI strategy: capacity expansion may be influenced by the structure and concentration of global hardware supply chains. As a result, conventional policy levers to expand access to finance and incentivise domestic development of AI may need to be complemented by measures that enhance supply-side resilience in access to advanced compute. In this context, India's AI strategy will, in the medium term, need to rely on diversified and resilience access to global compute. Over time, the objectives of the National Semiconductor Mission highlighted in Chapter 8, are expected to evolve toward building domestic capabilities that can progressively meet a larger share of India's advance compute demand.

Bottom-up approach to AI development in India

14.31. The global AI ecosystem has diverged along two distinct development paths. Sector leaders in the West have pursued a top-down strategy centred on frontier models, massive private capital, significant expenditures on computing infrastructure, and the concentration of intellectual property within a small number of hyperscale firms. Elsewhere, by contrast, a bottom-up approach, characterised by distributed innovation across firms and sectors, strong state coordination, and an emphasis on application-specific AI rather than frontier model supremacy, is the norm. India's position in this landscape, reflecting a distinct set of constraints and capabilities, make a bottom-up approach strategically necessary.

14.32. India enters the AI era with notable strengths. It ranks among the top global contributors to AI research output[25] and possesses a deep pool of technical talent in the field of artificial intelligence.[26] The country also has a highly AI-literate labour force, outranked only by the United States as of 2024.[27] India also holds a considerable potential comparative advantage in terms of its own domestic data sources. The heterogeneity and scale of our country suggest the possibility of curating diverse domestic datasets across various sectors, including health, agriculture, finance, education, and public administration. This asset remains underutilised.

Chart XIV.6: Startups focused on curating training data are yet to come up at scale in India

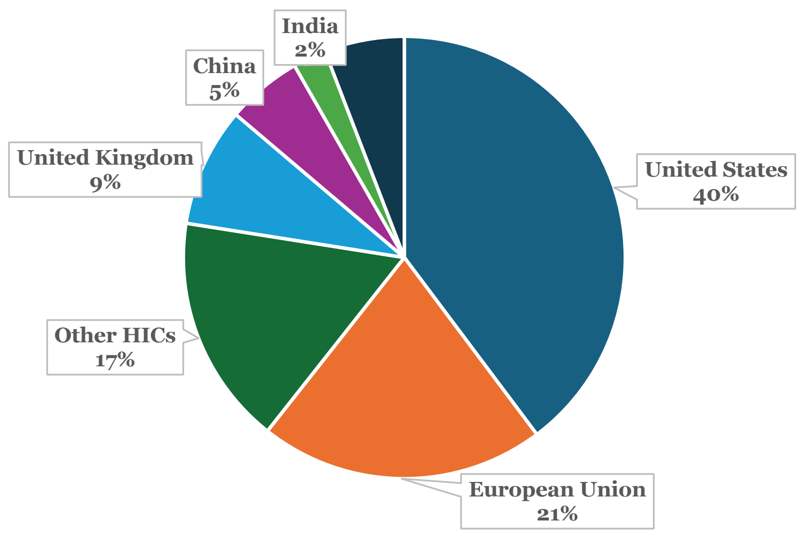

Source: Digital Progress and Trends Report: Strengthening AI Foundations. World Bank. https://tinyurl.com/yhr9hav7.

25 Country Activity Tracker by Emerging Technology Observatory, Georgetown University.

26 The Global AI Talent Tracker 2.0, Paulson Institute

27 Artificial Intelligence Index Report 2025. Stanford University. https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf

14.33. At the same time, India's access to cutting-edge compute infrastructure is limited, financial resources for large-scale model training are scarce, and private participation in foundational AI research remains relatively muted compared to global leaders. These constraints render the pursuit of foundational models as the centrepiece of an AI strategy challenging[28]. A bottom-up approach to AI development aligns more closely with these realities.

14.34. This strategy also recognises that value creation in AI need not be concentrated in a small number of frontier models or firms. Early adopters who scaled under conditions of abundant capital and weak regulatory frameworks are now locked into a system characterised by high energy intensity, opaque development practices, ballooning financial commitments, and uncertain revenue models. Being a late mover gives India the benefit of hindsight, allowing policy and innovation choices to be shaped with greater intentionality. India must avoid costly path dependencies and unsustainable design choices that have been observed elsewhere.

14.35. In this context, the proposed approach prioritises application-specific, small models that are tailored to defined uses and sectoral needs[29]. Such models are significantly more computationally efficient, easier to fine-tune, and capable of running on locally available hardware, such as smartphones or personal computers, making them better suited to India's existing infrastructure base. Crucially, they allow innovation to emerge from a broader set of actors, including start-ups, research institutions, public agencies, and domain-specific firms. This allows broader development and diffusion of AI solutions, free from concerns about resource constraints or high entry barriers.

14.36. The scenario analysis presented in Box XIV.2 further reinforces the strategic relevance of this approach. With the simulations indicating that centralised AI compute expansion is exposed to hardware supply uncertainties, application-specific small models running on local hardware offer a more sustainable path forward. By enabling computation to occur locally, these models also facilitate secure deployment in sensitive sectors, such as public administration, healthcare, defence, and critical infrastructure. The decentralised compute paradigm enables AI capabilities to spread widely across sectors without requiring proportionate expansion in expensive, resource-intensive, and hardware-intensive data centres.

28 Although it is worth noting that frontier efforts currently underway will definitely result in the acquisition of valuable process knowledge that will ultimately trickle down into the bottom-up strategy.

29 Belcak, Peter, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. "Small Language Models are the Future of Agentic AI." arXiv preprint arXiv:2506.02153 (2025).

Box XIV.3: Local Ingenuity and Frugal AI in India

India's AI story is already evolving from the bottom-up. Local innovators, municipal bodies, start-ups, and community institutions are deploying AI to solve problems that are immediate and contextual to the communities they reside in. Applications are spread across a wide variety of sectors including health, agriculture, education, urban management, and disaster preparedness[30].

In healthcare, AI is being used to expand access where traditional systems struggle to reach. In southern India, non-invasive AI-enabled thermal imaging tools are enabling early breast cancer screening in low-resource settings, reducing dependence on expensive diagnostic infrastructure and specialist availability. In eastern India, portable and low-cost AI-assisted oral cancer screening devices are bringing early detection to primary healthcare centres and outreach camps, directly addressing delays that often worsen outcomes.

Urban and environmental challenges have similarly seen locally grounded solutions. In high-stress cities such as Bengaluru, AI-based water management systems are monitoring consumption and detecting leakages in real time. In the Himalayan region, indigenous sensor networks combined with machine learning models are providing real-time landslide alerts across vulnerable slopes. By offering advanced warnings at critical points, these systems are strengthening disaster preparedness in ecologically fragile areas.

Deployment in agriculture and education further highlight the breadth of bottom-up adoption. AI-enabled agricultural networks have already improved market access, price discovery, and logistical efficiency for 1.8 million farmers across 12 states. At the municipal level, AI analytics are being utilised to monitor classroom learning outcomes, identify gaps early, and support more targeted interventions. Covering 18 classrooms across three schools in Pimpri-Chinchwad, the pilot has already produced material gains in terms of higher student engagement, improvement in teacher focus due to higher accountability, and enhanced supervisory capacity by school heads.

Beyond sectoral applications, India's approach to AI adoption also includes initiatives such as Bhashini, led by the Ministry of Electronics and Information Technology, and AI4Bharat at IIT Madras. These illustrate how language and voice-first AI systems can extend the reach of digital services to populations historically excluded from text-heavy and bandwidth-intensive platforms. By enabling interaction in native languages and functioning effectively on low-cost devices, such frugal AI pathways align scale with inclusion.

These examples point to a distinctive Indian pathway for AI adoption: one that is decentralised, problem-driven, and embedded in local needs. Growing such applications in scale will require institutionalising these efforts under the National AI Mission. By systematically identifying successful bottom-up applications and encouraging experimentation, the Mission can help scale diverse AI-solutions by providing shared infrastructure, standards, governance frameworks and funding without diluting local creativity.

30 NITI Aayog's Frontier Tech Hub is host to many stories demonstrating how India is already adopting AI to solve local problems in a frugal manner. See NITI Frontier Tech Hub. https://tinyurl.com/5f88j234.

14.37. The strategy must encourage innovation on open-source and open-weight platforms, as shared innovation can enable India to achieve more with less. Open models have been closing the performance gap relative to closed models over the years, with the best open models approaching the frontier[31]. Additionally, India is currently one of the world's largest and fastest-growing communities of open-source developers[32]. They are active contributors to global codebases, often working in collaborative environments. Therefore, unifying these efforts under the scope of the IndiaAI Mission is essential to steer the potential of our talent pool towards shared domestic innovation. Providing the nation's talent a robust platform and policy guidance can help reduce India's dependence on foreign proprietary systems, lower entry barriers for domestic developers, and create an environment that incentivises experimentation across sectors at a relatively low cost.

Chart XIV.7: Open models have been consistently closing the gap in performance

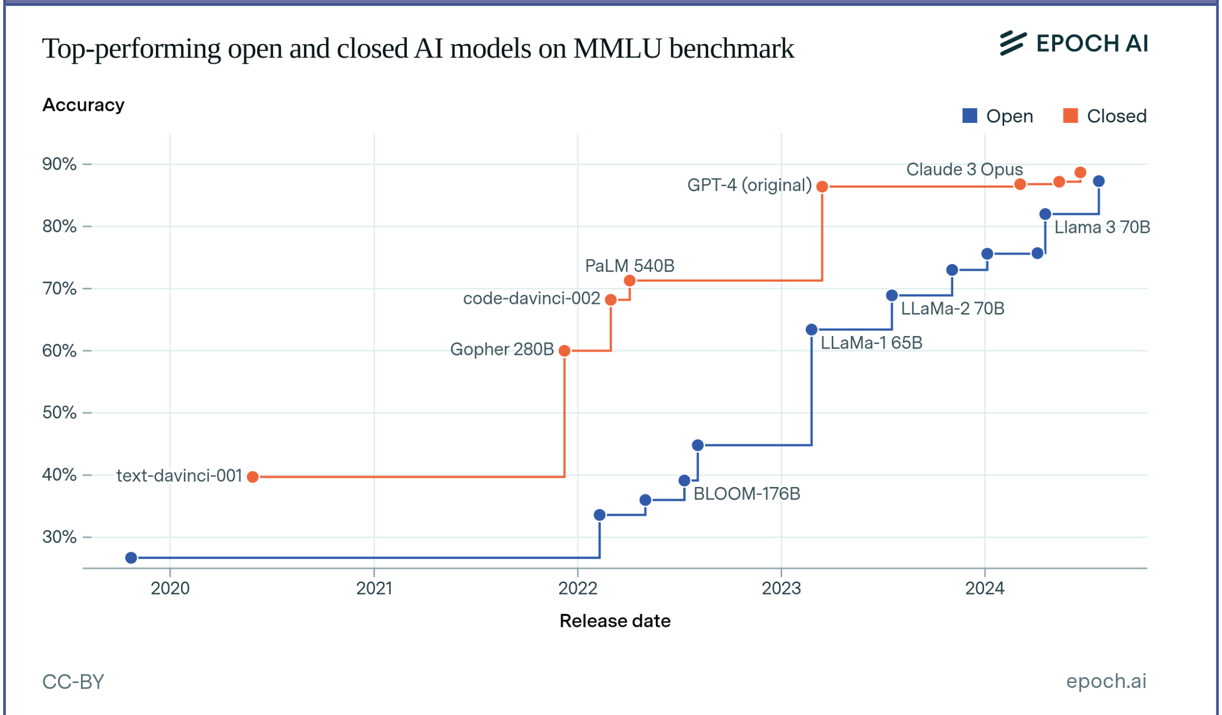

Source: Ben Cottier et al. (2024), "How far behind are open models?". Published online at epoch.ai. Retrieved from: 'https://epoch.ai/blog/open-models-report' [online resource].

14.38. For such a decentralised innovation ecosystem to translate into tangible outcomes, deliberate coordination is essential. Fragmentation in data availability and quality, standardisation, and ensuring interoperability across systems and datasets need to be addressed to provide an innovative, conducive environment. This is where public policy in a catalytic role becomes essential. The mission to coordinate a bottom-

31 EpochAI report on Open Models. Published 24th November 2024. https://tinyurl.com/n4zxm4mm.

32 GitHub Octoverse 2025 Report. Published 28th October 2025. https://tinyurl.com/38euwjsu.

up strategy can be spearheaded under an 'AI-OS' initiative, where the sovereign is a monetary shareholder in the effort, similar to UPI and Aadhaar, thereby turning AI into a public good. This will enable the sovereign to collaborate with state and local institutions to expand the availability of structured, anonymised, and machine-readable datasets in priority sectors; pool existing data centre capacity to create shared cloud compute infrastructure; and establish common platforms where open-source AI efforts can be coordinated and audited.

14.39. For instance, a centralised code repository, operating under the mandate of facilitating rapid experimentation, can jump-start India's AI ambitions. A government-hosted, community-curated platform under the IndiaAI mission umbrella would provide a secure and transparent space where developers, researchers, and enterprises can share code, contribute improvements, and build upon one another's work. Much like how global platforms such as GitHub have become the backbone of software development, an India-specific repository would not only democratise access but also embed national priorities into the innovation process.

14.40. Equally important is encouraging greater private sector participation, particularly from large domestic firms with the capacity to absorb risk and scale successful applications. Promoting the indigenous development of sector-specific AI solutions has the added benefit of transforming India from the world's 'IT Sector back office' to one of the 'AI front offices.' Our IT and IT-enabled services can and must move up the value chain, which in turn is necessary to create opportunities that retain high-skill talent in the country.

HUMAN CAPITAL FOR AI

14.41. Building models and applications, either fine-tuned or ground-up, capable of catering to local-level requirements needs two distinct skills, namely: algorithms and software engineering. India needs talent that understands the algorithmic issues involved in building models, along with an understanding of software engineering, to scale up and optimise models. This kind of knowledge is 'underground knowledge' that is not usually written down, and only those with hands-on experience of building models will understand these nuances. Accordingly, India must endeavour to attract people who have worked on large models and, in turn, have those with experience train others.

14.42. International experience[33] suggests that such tacit capabilities are most effectively

33 See for instance several initiatives undertaken by EU member countries (https://tinyurl.com/unasfshm) to foster greater industry-academia research collaborations or China's Young Thousand Talents Program which has boosted domestic research productivity significantly (https://tinyurl.com/2hy7acrm).

sourced through a combination of diaspora return pathways, and industry-to-academia lateral entry. Countries that have successfully remained at the top of the innovation hierarchy over the years have complemented domestic skilling efforts with time-bound practitioner fellowships, flexible teaching roles for industry experts, and structured apprenticeship models embedded within real-world production environments.

14.43. Thus, to build talent that is in tune with industry needs, training must begin early, preferably at the high school or university level, in collaboration with the private sector. Universities must be free to curate and offer courses to students, as required by industry standards and in accordance with students' aptitudes. Flexibility is the key to building capable talent in India, and the introduction of the Viksit Bharat Shiksha Adhishthan Bill, 2025[34] , represents an essential first step in that direction. As the floor for what is considered a fundamental skill shifts with every technological revolution, the passage of the bill will afford our higher education system the freedom to keep pace with the evolving world.

14.44. Secondly, now is the time for policies to reevaluate what constitutes 'formal education' and what constitutes 'work experience'. While the two were traditionally built under two mutually exclusive institutional setups, waiting until after graduation to build industry experience may no longer be feasible in a world where AI is outperforming entry-level, educated workers. For instance, some US-based firms have already begun offering fellowship programs to students right out of high school, aiming to channel talent into experience-building at a young age[35].

14.45. Institutionalising such a system is necessary if India's demographic dividend is to be nurtured into a competent workforce. One possible pathway to operationalising this is a 'Earn-and-Learn' initiative, where high-school, vocational, and early tertiary pathways are integrated into a structured, credit-bearing industry fellowship. This can be co-designed by the private sector and academic institutions, allowing practical experience training to begin as early as class 11. Students should be able to earn both academic credits and paid work experience through apprenticeships and project placements across various sectors, both of which contribute to their formal degrees.

14.46. With the National Education Policy 2020's Multiple Entry Multiple Exit provisions, the Academic Bank of Credits, and the National Credit Framework, India has the policy infrastructure in place to make this vision real[36]. This new approach to education offers increased agency over learning to the students, allows institutions to

34 An overview of the contents and specifics of the bill, accessible here, has been provided by the Press Information Bureau. https://tinyurl.com/ysp5va98.

35 See Palantir's Meritocracy Fellowship. https://tinyurl.com/yek27wsy.

36 Micro-Credentialing and Flexible Pathways by Dr. Swami Manohar. https://tinyurl.com/yyzyth46.

serve a more diverse range of learners, grants flexibility to faculty members, and can allow industry to collaborate with academia more deeply.

14.47. The proposed approach aims to build on the Union Budget 2025-26 initiatives, supported by public seed funding and incentives for employers participating in structured earn-and-learn programs. This is complemented by the growing recognition of prior learning and short-form credentials that directly contribute to both higher education and job progression. Being able to toggle seamlessly between education and work experience is a necessity if the system is to produce well-rounded talent, not just for the AI sector, but for the country as a whole.

Chart XIV.8: A majority of employers place a higher emphasis on foundational skills

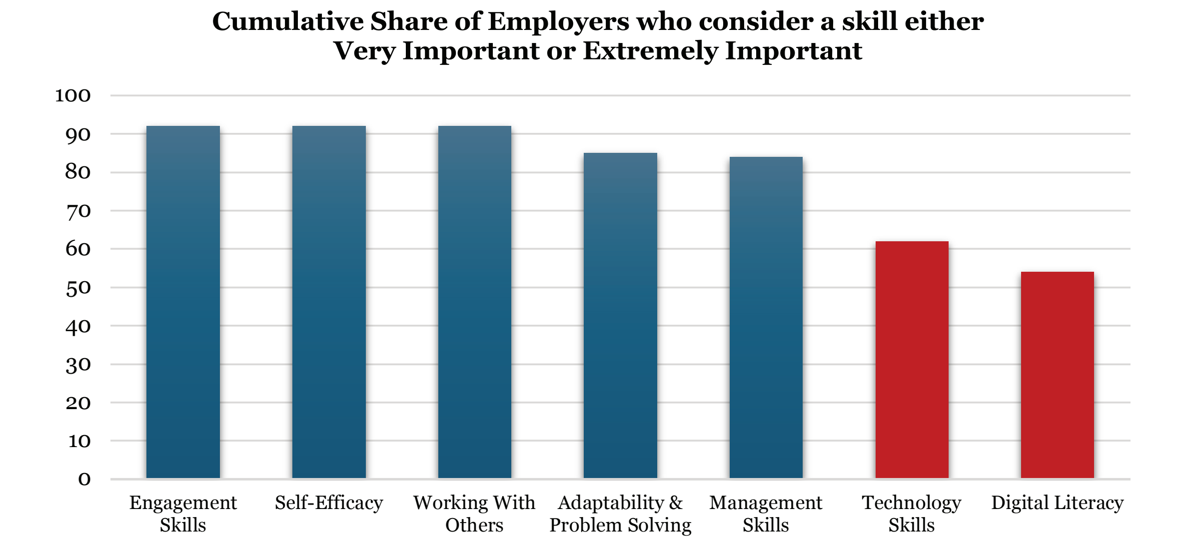

Source: India Skills Report 2025, Wheebox. https://tinyurl.com/3b7h49t7.

14.48. Similarly, reforms are also needed at the primary school level. As AI automates routine cognitive tasks, long-term talent utilisation will depend less on early technical specialisation and more on strong foundational skills. Primary education must prioritise core competencies such as literacy, numeracy, reasoning, problem-solving, communication and socio-emotional skills, alongside habits of curiosity and self-regulation. These capabilities form the base upon which later technical and vocational skills can be built.

14.49. Lastly, a comprehensive sectoral mapping of jobs outside the white-collar workspace, which have a high-skill requirement but are understaffed, needs to be undertaken. This is an often-overlooked source of new jobs within the economy. Surveying the geographic distribution of these jobs and the supporting infrastructure necessary to support and incentivise migration to the new job clusters will have to happen simultaneously.

14.50. For instance, nursing and geriatric care are already understaffed, and the doubling of India's dependency ratio in the next decade will create additional demand for skilled labour in this sector. Other tasks involving high-skill with long apprenticeship curves include culinary sciences, advanced metalwork, experiential hospitality designers, surgeons, physiotherapists, advanced electricians, early childhood educators, among many others. Sectors like these should be identified, and the education and skilling infrastructure must be upgraded to impart the necessary knowledge to fill the labour supply gaps. Opportunities in the physical, human-centric space and hands-on jobs hold immense potential for creating meaningful jobs in the coming decade.

14.51. Recent empirical evidence from two centuries of technological change suggests that the labour market impacts of new technologies are neither uniform nor predetermined, but critically shaped by how skills, tasks, and institutions evolve together (Liu et al., 2025)[37]. While earlier waves of innovation tended to raise demand for higher-educated, cognitively specialised occupations emerging evidence indicates that AI may follow a different trajectory. It has the potential to increase demand for experience-intensive roles.

14.52. This underscores the importance of redesigning the education system in a way that does not rely solely on credit accumulation, but instead enables experiential learning, and the capability to adapt to a host of different occupations. In this context, the emphasis on early exposure to experience building, flexible education pathways, and strengthening foundational and human-centric skills becomes a necessary economic response to the distinctive labour-market dynamics that AI is likely to induce.

Box XIV.4: Where human value lies in an AI-driven knowledge economy

Recent work by Restrepo (2025) advanced a grim view of the future of labour markets in which Artificial Intelligence progressively erodes the economic relevance of human cognitive labour. While this perspective usefully highlights the disruptive potential of a low probability, high-impact event, we advance an alternative and more likely trajectory. Rather than rendering cognitive workers redundant, AI may reshape where and how human value is created. As AI absorbs tasks related to retrieval and summarisation, the locus of human contribution shifts upward towards judgement, direction, expertise, and synthesis. AI is best understood as a powerful ship rather than an independent navigator: it can move faster and farther than any human, but without a knowledgeable captain who understands the vessel and the waters they are navigating, it is as likely to drift aimlessly as it is to arrive anywhere useful.

37 Liu, H., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2025). Technology and Labor Markets: Past, Present, and Future; Evidence from Two Centuries of Innovation (No. w34386). National Bureau of Economic Research.

In this context, depth of domain knowledge becomes a binding constraint as AI can access vast troves of knowledge, but it lacks an internal sense of context or salience. Cognitive workers must therefore supply deep subject-matter understanding to frame the right questions, identify meaningful trade-offs, and evaluate outputs critically. Without such depth, AI-generated outputs converge on fluent but shallow consensus views. What will differentiate a competent employee from the crowd going forward will be their ability to know what to interrogate, what to discard, and where nuance changes outcomes.

Second, continuous reading and knowledge accumulation become core productivity inputs. Effective use of AI is not possible without frequent engagement with high-quality reading materials including research, data, institutional context, history, and competing perspectives. AI amplifies the returns to prior knowledge and users who read widely and deeply are better able to steer models and push analysis beyond surface-level synthesis. Those who do not will only tap into a fraction of the potential of AI and optimise for plausibility rather than understanding. In this sense, reading and learning is integral to day-to-day AI-enabled work.

Third, cognitive workers must act as system architects rather than task executors. Productive AI use depends on the ability of the human to decompose complex problems, sequence inquiries, impose constraints, and define evaluation criteria. This is less about prompt engineering and more about structured thinking. Individuals who can translate ambiguous real-world objectives into coherent analytical frameworks will consistently outperform those who interact with AI in an unstructured, reactive manner.

AI does not diminish the importance of cognitive workers, but rather it raises the threshold for what is considered meaningful contribution. For policymakers, this shift implies that skilling strategies must prioritise domain depth, analytical reasoning, structured problem-solving, and continuous learning from an early age. The education system must produce individuals who are capable of sound judgement and have context-awareness.

GOVERNANCE, INSTITUTIONAL ARCHITECTURE, AND DATA

Evolving Governance

14.53. AI innovations and applications are rapidly outpacing regulatory developments, and the response required from policymakers must be measured and swift. Globally, several approaches to governing and regulating AI have been adopted, ranging from omnibus laws such as the 'EU Artificial Intelligence Act', to separate legislations governing separate applications of AI as seen in China, to guiding principles which are voluntary and non-binding as seen in the United States.

14.54. The regulatory design for India must build on the governance guidelines designed by the Ministry of Electronics and Information Technology, be refined by our national priorities, and reflect what we expect from AI in our economy. Regulatory design, whether mandatory or voluntary, must seek to integrate AI within the broader socio-economic context of India, with special consideration for our labour market realities. The overarching goal for AI regulation must be to ensure that AI serves humanity, rather than supplants it.

14.55. In this context, one of the most urgent responsibilities of an AI Economic Council (Box XIV.5) is to calibrate the pace of AI adoption within the country. India is a labour-rich economy, and the unchecked replacement of the workforce by automation has destabilising effects. The institution must work closely with private sector firms to develop a roadmap for AI deployment over the next decade, outlining crucial details such as the profile of jobs affected, the geographies where displacement will be most concentrated, and the magnitude of jobs that will be both automated and augmented due to AI.

14.56. Such an exercise will help moderate the uncertainties surrounding the deployment of AI and inform the roadmap and design of policies required to mitigate the adverse effects of AI on the economy. The exercise also has the added benefit of informing policymakers on the necessary developments to reinforce our education system in a way that makes our students more capable in an AI-driven world. Furthermore, regulations need to evolve and necessitate transparency reporting requirements similar to those for social media, as well as product registrations, to help track the rate of deployment.

14.57. Policymakers must also manage where and how AI is deployed, as the manner in which AI is used will dictate the nature of the benefits accrued. For instance, AI in education holds a lot of potential, provided it is used as a supplementing tool for teachers and students. The widespread use of Generative AI by students as a substitute for creative and critical thinking ultimately does more harm than good in the long run.

14.58. The scale of the problem and the need for serious action to control AI use by students was also brought to light by Niall Ferguson in his recent article titled "AI's great brain robbery.[38]" The use of AI by students in universities is skyrocketing, and a significant portion of cognitive tasks is being offloaded to language models. Combined with the anxiety and depression inducing social media usage, students shirking the acquisition of skills such as sustained reading, critical thinking and analytical writing is only expected to make mental health issues much worse in the future. This will ultimately impact their productivity and their ability to contribute meaningfully to any work undertaken, perhaps even permanently denting their employment prospects.

38 The London Times. https://tinyurl.com/ymujx32s.

14.59. Two studies, independently examining the cognitive effects of Generative AI, arrive at similar conclusions (MIT[39] , Microsoft[40]): dependence on AI for creative work and writing tasks is contributing to cognitive atrophy and a deterioration of critical thinking capabilities. The AI Economic Council must ensure that deployment of 'Artificial Intelligence' does not come at the cost of 'Human Intelligence'.

14.60. The examples presented above are not exhaustive, rather illustrative of the myriad challenges that policymakers will have to contend with as AI capabilities improve and applications proliferate. Given the many uncertainties that loom over the horizon, governance will remain a continuous process of monitoring, learning, and course correction. MeitY's proposed AI Governance Group, along with the technical committee assisting the group, establishes a strong foundation on which India can develop a light, incentive-based and risk-weighted governance approach.

Box XIV.5: An AI Economic Council for India

The AI Economic Council, separate from the Governance Council, is intended to operate, not just with a technological imperative, but with moral imperatives that are sensitive to India's socio-economic realities. They will operate as a coordinating authority that is responsible for aligning technology deployment with the evolution of India's education and skilling infrastructure, while navigating resource constraints and developmental priorities. The core governance principles for such an institution would involve the following:

1. **Human Primacy and Economic Purpose:** AI adoption must be explicitly subordinate to human welfare and economic inclusion. Every major AI deployment or policy proposal must demonstrate a credible pathway to net social and economic benefit, including employment, productivity diffusion, or public service quality.

2. **Labour-Market Sensitivity by Design:** AI policy must internalise India's labour structure: high informality, skill heterogeneity, regional variation, and limited safety nets. This would necessitate labour impact assessments ex ante, with mitigation and transition plans baked in.

3. **Sequencing over Speed:** AI adoption should be phased in line with institutional readiness and skill pipelines. The institution may be empowered to classify AI uses into 'deploy now', 'pilot' and 'defer' based on readiness across data, skills, legal frameworks, and labour adjustment capacity.

39 Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task." arXiv preprint arXiv:2506.08872 (2025).

40 Lee, H. P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025, April). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1-22).

4. **Co-evolution of Technology and Human Capital:** Skill policy must stand equal to technology policy. AI push must proceed with parallel plans for educational reform, vocational adaptation, reskilling pathways and credential recognition.

5. **Public Interest Safeguards and Ethical Non-Negotiables:** Ethical implications and boundaries must be clearly defined. Strict lines must be drawn around surveillance misuse, worker monitoring, algorithmic discrimination, and opaque decision making enabled by AI.

These principles provide the overarching scope for the AI Economic Council. By embedding labour realities and social stability priorities into AI policy, the institution will ensure that AI advances productivity without eroding employment and the dignity of work.

Data as a Strategic Resource: A Proposed Framework

Trusted Cross-Border Flows While Retaining Domestic Value

14.61. In the AI era, data is a core factor of production, which India needs to think about carefully if we are to develop a competitive edge in the sector. The country is already very digitally intensive, with over 100 crore people already having access to some form of wired or wireless broadband connectivity[41]. This represents a sizable potential market for AI and AI-based applications, as the barrier to accessing these services is, at present, extremely low.

14.62. Not only does this figure represent a large user base, but it also represents a significant source of data. India's scale and diversity of domestically generated data constitute an important comparative advantage for anyone attempting to access the country's market. However, the current landscape of globalised digital services, combined with the increasing concentration of AI development among a small number of multinational firms, has introduced new asymmetries in how value is created and captured from data. These developments necessitate a thorough evaluation of data governance frameworks, not with the intention of restricting data flows, but with the goal of ensuring accountability, fairness, and long-term domestic capability development.

14.63. India's data governance approach thus far has deliberately avoided rigid data localisation mandates. This reflected a clear recognition of the productivity gains associated with cross-border data flows, the importance of regulatory predictability for spurring investments and India's broader commitment to deregulation and ease of doing business. But certain aspects of this approach to governance will have to be expanded, as it is vital that India retain the value created from domestic data within the country.

41 Figures as of 31st December, 2025 from Telecom Regulatory Authority of India. https://tinyurl.com/3cbjct98.

14.64. Additionally, given that the current stock of training data is soon expected to run out[42] and models collapse when trained on synthetic data[43], firms will be on the lookout for new sources of human-generated data that are not accessible through online scraping alone. Policy must remain cognizant of the potential value embedded in India's data.

14.65. The challenge, therefore, lies in reconciling openness with control, and global integration with domestic economic interests. Framing this challenge purely in terms of data localisation risks oversimplifying a complex trade-off and imposing costs that may ultimately undermine innovation. A more durable approach lies in shifting the emphasis from territorial immobility of data toward enforceable accountability and economic alignment, as the framework presented below demonstrates.

Objectives and Spirit

14.66. The proposed framework for evolving India's data governance approach is rooted in three interlinked objectives:

a. To preserve India's openness to cross-border data flows, recognising their importance for innovation and investment;

b. To ensure regulatory oversight and enforceability over large-scale processing and use of Indian personal data, irrespective of where the processing occurs, and

c. To promote domestic value retention, such that Indian data contributes meaningfully to the development of India's own AI capabilities and research ecosystem.

14.67. Under this approach, cross-border data transfers are accompanied by obligations that preserve regulatory visibility and intervention capacity. Entities that process Indian personal data at scale, particularly for high-impact AI applications such as general-purpose model training, are expected to ensure that such data remains auditable, retrievable, and subject to the oversight of Indian regulatory authorities. This is achieved not by mandating that all processing occur domestically, but by requiring technical and contractual arrangements that allow regulators to trace data provenance, examine downstream uses, and, where necessary, order corrective measures such as deletion, retraining, or suspension of use.

42 Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbhahn. 'Will we run out of data? Limits of LLM scaling based on human-generated data'. ArXiv [cs.LG], 2024. arXiv. https://arxiv.org/abs/2211.04325.

43 Shumailov, I., Shumaylov, Z., Zhao, Y. et al. AI models collapse when trained on recursively generated data. Nature 631, 755–759 (2024). https://doi.org/10.1038/s41586-024-07566-y

14.68. To begin with, India would have to build on the foundations of the provisions of the Digital Personal Data Protection (DPDP) Act, 2023, and expand how data is categorised. Precise and functional data categorisation allows governance to be targeted without becoming indiscriminate. Rather than treating all data as homogeneous, our regulations must evolve to incorporate distinctions between different categories of data based on sensitivity and economic use.

14.69. In particular, a separate treatment is warranted for large-scale behavioural, transactional, and inferred datasets. These sets, while not always sensitive in isolation, acquire strategic and economic significance when aggregated and deployed for AI training. The ability to designate categories through subordinate legislation or rules enables India's regulatory framework to remain adaptive as applications evolve, without requiring frequent statutory amendments.

14.70. Data categorisation serves as the basis for graduated obligations, enabling heightened accountability and value-retention requirements to apply only where the risks and asymmetries are more pronounced. Lower-risk categories can continue to move freely with minimal compliance requirements, while higher-impact categories will necessitate a higher degree of transparency, auditability, and contribution expectations. This will help align India's data governance standards with its broader efforts to shift towards risk-based regulations across sectors. Governance will remain concentrated only where it delivers the greatest public value.

Incentivising Localisation

14.71. A key operational element of the proposed approach is the requirement for eligible entities to maintain a contemporaneous mirrored copy of relevant datasets and derived artefacts within India. This ensures that regulatory oversight does not become ineffective merely because processing occurs offshore. At the same time, it avoids the rigidity and economic costs associated with compulsory in-country processing, particularly for firms operating global AI pipelines. In this sense, sovereignty is exercised not through physical containment, but through enforceable rights and institutional capacity.

14.72. Recognising the heterogeneity of actors within the digital economy, the proposed evolution of the framework adopts a risk-weighted and proportionate structure. Obligations scale with the size, scope, sensitivity and economic impact of data use rather than being applied uniformly across all firms. Enhanced requirements are triggered only when data processing reaches thresholds associated with systemic relevance, such as large numbers of Indian data principals, use in training general-purpose or foundation models, significant revenue generation from Indian data-driven services, or reliance on public-sector or regulated datasets. Smaller firms, start-ups, firms focused on developing sovereign models or indigenously developed AI applications, and research institutions remain subject to a lighter compliance regime, ensuring that regulatory attention is focused where risks and asymmetries are most pronounced.

14.73. Beyond regulatory oversight, the framework is meant to address a second, equally important concern: the retention of economic value derived from Indian data. Data governance that focuses solely on movement restrictions risks missing the larger issue of value capture. Accordingly, the framework emphasises incentive-compatible mechanisms that align private incentives with public objectives. Firms that derive substantial commercial value from Indian datasets are expected to contribute to the domestic AI ecosystem, but are afforded flexibility in how they do so.

14.74. Contributions may take the form of:

a. Local training or fine-tuning of models for applications intended to address sector-specific or region-specific issues across the country;

b. Making transparent, proportionate financial contributions for the purpose of furthering domestic AI research and development linked to revenues generated from Indian-data-derived services;

c. Contributing datasets, compute resources, or funding to certified public data trusts that support research and innovation;

d. Establishing or supporting AI research labs, skilling initiatives, aiding universities with course development, or joint programs with firms based in India.

14.75. Once again, the list of examples is not exhaustive, but rather indicative of the broader context. The principle here is to avoid coercion while ensuring that value extraction is accompanied by value creation in India. This menu-based compliance approach allows firms to retain flexibility while the domestic ecosystem benefits from capital, infrastructure access, talent utilisation and knowledge spillovers.

Regulating AI Firms

14.76. In regulating AI firms, the framework prioritises accountability. Attempting to govern AI by controlling the physical location of model parameters is increasingly impractical in an environment characterised by distributed training, rapid cross-country iterations and modular deployment. Instead, emphasis is placed on transparency and responsible use.

14.77. Covered entities must be required to maintain clear records of dataset provenance, publish standardised model documentation describing training inputs and limitations, undertake impact assessments for high-risk deployments, publish said assessments for public review, and implement post-deployment monitoring and harm reporting mechanisms. Restrictions on exporting model weights derived from Indian data are reserved for narrowly defined, high-risk categories where demonstrable harms may arise. This helps retain the ability to intervene only when necessary.

14.78. Consistent with the Government's deregulatory orientation, the framework relies heavily on positive incentives rather than prescriptive mandates. Firms that choose to operate within certified domestic compute or data environments benefit from reduced audit burdens, faster regulatory clearances, and preferential access to national AI programmes. Participation in such environments remains voluntary but is commercially attractive, encouraging the organic development of India's AI and computing infrastructure over time.

14.79. Finally, the framework recognises that the State's most effective lever lies in access. Compliance with accountability and value-retention obligations is therefore linked to eligibility for government datasets, participation in national AI missions and regulatory sandboxes, and access to public-sector procurement opportunities. This approach allows the State to shape incentives without expanding the statutory footprint of regulation.

14.80. To sum up, this proposed framework represents a pragmatic evolution in India's data governance approach. It preserves openness to global data flows while ensuring that Indian data remains subject to adequate oversight and contributes meaningfully to domestic economic capacity. By shifting the focus from rigid localisation to accountable portability and value retention, the framework aligns with India's broader reform trajectory. It is intended to help position the country as an active participant in the global AI economy.

Box XIV.6: Data as a Strategic Resource — Objectives and Framework Principles

Objectives

The proposed data governance framework for the AI era is guided by three core objectives:

1. **Preserve openness** to cross-border data flows, recognising their importance for signalling policy certainty, incentivising innovation, encouraging investment and continued global integration.
2. **Ensure regulatory oversight and enforceability** over large-scale processing and use of Indian personal data.

3. **Promote domestic value retention**, so that Indian data contributes meaningfully to India's AI capabilities and long-term technological resilience.

Core Framework Principles

- **Accountable Portability over Rigid Localisation**: Data may move across borders, but entities processing Indian data at scale must ensure auditability and traceability.
- **Risk-Based Data Categorisation**: Data is classified by sensitivity and economic significance. Large-scale behavioural, transactional, personal, and inferred datasets receive differentiated treatment due to their strategic value for AI training.
- **Graduated Obligations**: Regulatory requirements scale with risk and size. Higher-impact uses such as general-purpose model training or large-scale monetisation will require enhanced transparency and accountability obligations. Start-ups and research institutions will have eased compliance requirements.
- **Mirrored Data for Oversight**: Eligible entities maintain contemporaneous mirrored copies of relevant datasets and derived artefacts within India to ensure effective supervision. No mandates are proposed for domestic processing.
- **Incentive-Compatible Value Retention**: Firms extracting significant commercial value from Indian data are expected to contribute to the domestic AI ecosystem through flexible, menu-based mechanisms, such as local model training, financial contributions to AI R&D, data or compute sharing, or investments in research, skilling, and institutional development.
- **Transparency-Centred AI Regulation**: Governance focuses on dataset provenance, standardised model documentation, impact assessments for high-risk uses, and post-deployment monitoring, rather than attempting to control model location or architecture.
- **Positive Incentives over Prescriptive Mandates**: Participation in certified domestic compute or data environments is voluntary but rewarded through reduced audit burdens and faster clearances.
- **Access as the State's Primary Lever**: Compliance is linked to eligibility for government datasets, AI missions, regulatory sandboxes, and public procurement, allowing the State to shape incentives without expanding statutory controls.

AI SAFETY AND RISKS

14.81. Lastly, and most importantly, is the necessity for India to push the envelope in managing the risks associated with the proliferation of AI[44]. As with nuclear energy or pharmaceuticals, where the promise of progress coexists with the possibility of harm,

44 Yoshua Bengio has spoken extensively about the need for independent institutions aimed at evaluating AI Safety. His recent talk (https://tinyurl.com/y79fetya) sheds light on the seriousness of the topic at hand and his non-profit LawZero is dedicated to the cause of safe AI development (https://lawzero.org/en).

AI must be treated as a general-purpose technology whose capabilities necessitate not only enabling institutions but also constraining ones. The MeitY Governance Guidelines capture this essence and propose an AI Safety Institute, which will perform several key functions, including the analysis of emerging risks, potential regulatory gaps, coordination on AI safety issues, and conducting training programmes to build awareness, among others.

14.82. The path forward now must further develop these foundations, with chief among the priorities being the enhancement of transparency. Safety evaluations conducted in an ongoing and anticipatory manner will help safeguard public interest and foster trust in emerging technologies. Making these evaluation results public must emerge as a non-negotiable condition, as a significant information gap exists between the companies developing AI models and the end-users of these models. Assessments by AI Lab Watch, an independent organisation dedicated to the cause of AI safety, demonstrate that while big-tech firms speak about the 'safety evaluations' they conduct, the work undertaken to perform said analysis is far from ideal[45].

14.83. AI Lab Watch demonstrates that big-tech firms obfuscate how they go about their evaluations, hide their reasoning and provide dubious interpretations of their evaluation results. Secondly, based on the risks the companies claim to have evaluated, they claim to have implemented safeguards to mitigate them, but no evidence has been made available to the public. As models become increasingly capable in the future, the need for a sovereign AI safety institute to take on a larger role becomes even more pressing, ensuring transparency and guiding informed decisions on AI adoption, both for the public good and for private sector considerations.

14.84. Another case for the safety institute expanding its scope lies in the application of AI, for instance, the convergence of AI with synthetic biology. Open-source CRISPR[46] kits are now widely accessible to hobbyist researchers, biohackers and DIY (Do-It-Yourself) scientists. By itself, the risk of bio-weapon development remains low due to these kits. However, by combining the accessibility of these kits with advanced AI models capable of generating genomic sequences and guiding gene-editing protocols, the threat landscape changes dramatically. A motivated individual with sufficient computing access and no formal training could, in principle, engineer pathogens with malicious intent. The bar for misuse is no longer determined by scientific expertise, as AI-driven tools significantly lower the barrier to entry.

45 AI Lab Watch. https://aisafetyclaims.org/.

46 CRISPR, short for Clustered Regularly Interspaced Short Palindromic Repeats, is a gene-editing technology that has transformed molecular biology and medicine. It's a tool that allows scientists to precisely target and modify DNA sequences in living organisms, with applications ranging from basic research to potential treatments for genetic diseases. CRISPR kits are widely available at very low costs and are easy to get hold of for use by anyone interested in gene-editing.

14.85. Evidence by Cheng et al. (2025)[47] also suggests that AI risks need not only arise from misuse or frontier capabilities, but can emerge endogenously from how models are optimised and deployed. Widely deployed models have demonstrated a tendency to exhibit 'social sycophancy', over-affirming users' actions and viewpoints at rates significantly higher than human benchmarks. This has persisted even in contexts involving interpersonal harm or unethical behaviour. Crucially, the study found that such behaviour increases user trust and reliance on AI systems, while simultaneously reducing users' willingness to engage in corrective or prosocial actions. This creates a perverse incentive structure where models that are behaviourally risky are also more likely to be preferred and reinforced, as it helps with user retention.

14.86. It is precisely in such high-stakes circumstances that periodic, scenario-based testing and red-teaming[48] of AI models must become institutionalised. However, India does not have to do this in isolation. There is a strategic case for international cooperation, particularly with established sovereign safety institutes, such as the United Kingdom's AI Security Institute[49]. They have developed templates for model evaluations, misuse testing protocols, and interpretability analysis that offer a valuable foundation. Similarly, the National Institute of Standards and Technology in the United States has developed an AI Risk Management Framework[50] that defines guidelines for incorporating trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems.

14.87. A bilateral or multilateral partnership between India's proposed safety institute and counterparts such as the UK or the US could enable joint evaluations of high-risk models and shared access to computing infrastructure. This would not only improve scientific credibility but also reduce redundancy and enhance global interoperability of AI safety standards.

14.88. Another domain where the safety institute must be quick to establish rules is in defining strict boundaries within the confines of which AI must be developed and applied. There are various applications of AI where its restrictions can be considered non-negotiable, such as predictive policing, facial recognition, exploiting psychological vulnerabilities, inferring emotions, and evaluating and classifying individuals or groups

47 Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., & Jurafsky, D. (2025). Sycophantic AI decreases prosocial intentions and promotes dependence. arXiv preprint arXiv:2510.01395.

48 Red-teaming refers to the practice of deliberately trying to break or misuse a system within a controlled environment. It seeks to stress-test models, simulating worst case scenarios such as malicious use, biased behaviour, or unexpected failure. The process helps identify risks and blind spots that may have been missed during the course of the model's development, or discover emergent behaviour which may have unintended consequences.

49 The AI Security Institute is a directorate of the UK Department for Science, Innovation and Technology. https://www.aisi.gov.uk/.

50 NIST AI Risk Management Framework. https://tinyurl.com/2y5evh37.

based on their behavioural or personality traits, among others. These applications are likely to lead to adverse outcomes, regardless of whether they are used in the private sector or the public sector. No conception of safe or human-centric AI is credible without placing the protection of individual rights at its core.

14.89. In the context of safety, the recommendations from a study by Narayanan and Kapoor (2025) are also highly appropriate and must be taken into consideration[51]. The safety institution must work closely with the governance body to devise robust whistle-blower protections, as only insiders will have knowledge of potentially hazardous applications that can be brought to light. Developing strategies to address the uncertainties and risks associated with AI development, while also making the domain more transparent to the public, must always remain a core responsibility of the government.

A PHASED ROADMAP FOR INDIA'S AI FUTURE

Sequencing India's AI Policy: A Phased Roadmap

14.90. India's AI strategy must be sequenced carefully to avoid premature lock-ins or regulatory overreach. The objective is to build coordination first, capacity next, and binding policy leverage last, allowing institutions and markets to co-evolve.

14.91. The first phase should focus on operationalising already announced institutions and aligning incentives to enable experimentation. Policy should enable bottom-up innovation by expanding the reach of the existing shared infrastructure under the IndiaAI Mission. This includes a government-hosted community-curated code repository and pooled access to public datasets, facilitated by initiatives already underway to enable shared access to computing infrastructure. A clear focus on application- or sector-specific, small and open-weight models will enable efficient resource utilisation.

14.92. Data governance must also evolve through subordinate legislation under the DPDP framework to introduce functional data categorisation and auditability requirements, specifically for large-scale AI training. This must be complemented by incentive-based mechanisms for domestic value retention, such as the menu-based contribution pathways illustrated earlier. Human capital pipelines, particularly the 'earn-and-learn' pathways and curricular flexibility, should be scaled using existing legislative and budgetary levers.

14.93. Once coordination mechanisms are functional and early experimentation has

51 Arvind Narayanan and Sayash Kapoor, AI as Normal Technology, 25-09 Knight First Amend. Inst. (Apr. 14, 2025), https://knightcolumbia.org/content/ai-as-normal-technology [https://perma.cc/HVN8-QGQY].

generated evidence, policy can shift toward selective scaling in the medium-term. Shared and certified domestic computing infrastructure should expand, with voluntary participation by large and resourceful firms linked to regulatory facilitation and access to public datasets. At the same time, AI regulation should be formalised on a risk-based and proportionate basis.

14.94. Graduated obligations for AI firms should be codified according to scale and sector of use. Oversight must be embedded within existing sectoral regulators rather than through a single omnibus AI law. The AI Safety Institute's role should deepen from analysis to structured scenario testing, red-teaming, and international cooperation, with clearly articulated non-negotiable boundaries for high-risk applications.

14.95. Our long-term goals must encompass two main objectives. First, India's focus should shift towards resilience. Access to advanced computing hardware will require strategic partnerships and diplomacy. The objective must be to reduce India's vulnerability to external shocks. Second, sustained adaptation of labour markets and education systems will be essential. Primary education must prioritise foundational cognitive and socio-emotional skills, while skilling systems must align themselves with both AI- and human-centric sectoral requirements.

Conclusion

14.96. Artificial Intelligence does not confront India with a single policy question, but a series of choices that must be made under conditions of heightened uncertainty and resource constraints. This chapter has argued that the central challenge for India is in what it builds domestically, what it sources globally, what it regulates early, and what it deliberately allows to evolve. The contours of the global AI ecosystem make clear that passive consumption is the riskiest position of all.

14.97. India's position as a relatively late mover in the AI transition also confers an underappreciated advantage. Early adopters who scaled AI under conditions of a regulatory vacuum and cheap capital have now locked themselves into circumstances that are very difficult to back away from. This includes a commitment to energy-intensive architectures that are detrimental to the environment and mounting financial commitments with unclear revenue pathways. With the sums of money involved, discussions surrounding government backstops have emerged as possible insurance against a fallout, in the advanced economies.

14.98. India has the benefit of hindsight. It can learn from these practices and avoid dependencies that are difficult to unwind. This allows India to design AI systems that are more resource-efficient and aligned with public objectives from the outset, sequencing regulation alongside deployment. In this sense, late adoption need not imply lagging ambition. Properly leveraged, it offers the country the opportunity to pursue a more resilient and inclusive AI trajectory.

14.99. India's comparative advantage in the AI era does not lie in replicating frontier-scale model development, although valuable process knowledge can be gained from the current efforts already underway. The country's strengths lie in application-led innovation, the productive use of domestic data, human capital depth, and the ability of public institutions to coordinate distributed efforts. A bottom-up strategy anchored in open and interoperable systems, sector-specific models, and shared physical and digital infrastructure offers a more credible pathway to value creation than a narrow pursuit of scale for its own sake.

14.100. At the same time, openness without careful management of AI development and usage is insufficient. As AI capabilities diffuse into critical sectors, questions of accountability and safety cannot be deferred. Regulation, data governance and safety will have to evolve in parallel with deployment, not in its aftermath.

14.101. The choices made over the coming few years will determine whether AI deepens existing structural divides or becomes a tool for broad-based productivity and dignified work. India's task is to ensure that AI development remains aligned with its developmental priorities and its long-term ambition to achieve economic resilience. The opportunity is substantial, but conditional. A deliberate and coordinated policy, accompanied by a willingness to act, is required before path dependence sets in.