



TAKSHASHILA
INSTITUTION

Why and How India Should Accelerate AI Diffusion

Bharath Reddy, Rijesh Panicker

Takshashila Discussion Document 2026-30
Version 1.0, September 2026

The paper examines why India should accelerate AI diffusion and proposes policy pathways for doing so.

Recommended Citation:

Bharath Reddy, Rijesh Panicker, "Why and How India Should Accelerate AI Diffusion", Takshashila Discussion Document 2026-30, Version 1.0, September 2026, The Takshashila Institution

©The Takshashila Institution, 2025

Contents

1	Executive Summary	3
2	Why Diffusion Matters	4
3	Why India Should Accelerate AI DiffusionA	5
3.1	Productivity of Indian firms	5
3.2	Public service delivery	6
3.3	India has substantial capabilities with which to pursue this opportunity	6
3.4	But, rapid AI diffusion also carries significant risks . .	7
4	The Limitations Of The No Intervention Path	7
5	The Four Conditions for AI Diffusion	9
5.1	Reliability at the task	10
5.2	Access to the service	11
5.3	Complementary capital	12
5.4	Permissions	13
5.5	Mapping the conditions to the AI stack	15
6	What India Should Do To Drive AI Diffusion?	16
6.1	Reliability at Task	17
6.1.1	Establish common evaluation standards and deployment-specific validation	17
6.1.2	Require evidence of performance across regional languages and user groups	17
6.1.3	Tie procurement and expansion to evidence of useful outcomes	18
6.1.4	Treat reliability as a continuing obligation . . .	18
6.2	Access to the Service	18
6.2.1	Make open-weight models a default option in public procurement and ensure interoperability and portability	18
6.2.2	Support an international consortium to sustain competitive open models	19
6.2.3	Pool public-sector inference demand while preserving choice and accountability	19
6.2.4	Reduce the friction in buying technology from abroad	20
6.2.5	Make underserved AI applications affordable through existing service networks	20
6.2.6	Treat electricity distribution reform as AI strategy	20
6.3	Complementary capital	21
6.3.1	Build practical implementation capabilities and budget for organisational change	21
6.3.2	Invest in the skills and judgement needed for AI adoption	21
6.3.3	Invest in open, common standards to enable vibrant AI markets	22
6.4	Permissions	22

6.4.1	Build institutional trust and give officials room to learn	22
6.4.2	Make data-protection requirements easier to apply in practice for AI use cases such as voice	22
6.4.3	Strengthen sectoral regulation with dedicated expertise and funding	23
7	Conclusion	23

1 Executive Summary

AI diffusion is the widespread adoption of AI technologies across firms, public institutions and sectors of the economy. The case for accelerating AI diffusion in India rests on three arguments. First, AI can improve firm productivity by making skills in activities such as marketing, software development, and inventory management more accessible to enterprises. Adoption can also help Indian firms remain competitive as AI changes the delivery of tradable services. Second, AI can improve public service delivery by helping frontline officials deliver services and enabling institutions to process information and coordinate work more effectively. Third, India's IT services industry and digital public infrastructure provide capabilities for integrating AI into operational systems. Domestic diffusion could improve services for Indian users while developing implementation expertise and products relevant to other countries.

Diffusion will occur without additional state intervention, but commercial incentives will favour users with purchasing power and organisations already equipped to adopt AI. Falling model prices will not by themselves supply representative data, reliable interfaces in underserved languages, integration into public institutions, trusted intermediaries or effective recourse. Public policy should address these gaps where benefits extend beyond individual providers, commercial demand is insufficient, or the government has responsibility for delivering the service.

The paper identifies four conditions for beneficial and sustained diffusion and suggests policy interventions to address the gaps identified:

- **Reliability at the task:** Systems must perform adequately under actual operating conditions. The government should fund reusable evaluation infrastructure and representative datasets, require deployment-specific validation, and link procurement and expansion to demonstrated outcomes. Evaluation must continue after deployment.
- **Access to the service:** Intended users must be able to obtain and use services affordably and consistently. Policy should support suitable delivery channels, competitive inference supply and interoperable procurement. Open-weight models should be evaluated alongside proprietary services, with arrangements that preserve supplier choice and continuity.
- **Complementary capital:** Adopters need the data, skills, workflows and supporting resources to act on AI outputs. Investment should cover implementation support, organisational change and continuing operations, including the downstream services needed to realise benefits.
- **Permissions:** Deployments need clear legal authority, institutional acceptance and accountability. Policy should clarify data-protection requirements, involve frontline workers and affected users, provide accessible correction and redress, and strengthen sectoral regulators with

This paper uses AI as an umbrella term to cover generative AI and other machine-learning applications. Its focus is on the adoption and use of these systems in economic activity and public services, and the conditions under which that use produces benefits.

AI use disclosure: ChatGPT was used for reviewing the paper, brainstorming logical inconsistencies and copyediting.

Nitin Pai and Pranay Kotasthane's inputs contributed to refining the ideas in this paper.

technical expertise and predictable funding.

2 Why Diffusion Matters

Diffusion is the widespread adoption of a technology across firms, sectors and populations. The claim that diffusion drives economic returns from general-purpose technologies (such as electricity, the Internet, or AI) is well established.

Jeffrey Ding makes the argument that durable economic benefits from general-purpose technologies depend not on pioneering it, but in having the skills and institutions to spread it widely. Mazarr's RAND monograph reaches a similar conclusion from a national power lens¹, concluding that "the competitive challenge of AI is primarily social, not technological". A strategy built only on dominance of the technology stack (models, chips, and data centres) will not lead to competitive advantage if societal integration is not prioritised.

Paul David's account of electrification illustrates the importance of complementary changes in enabling diffusion². Electric motors were commercially available in the 1880s, yet US manufacturing productivity did not visibly respond until the 1920s. The lag was due to organisational reform required to internalise the benefits through slow experience-based learning. But investments in complements might take a while to reflect in the productivity statistics.

Brynjolfsson et al. argue that AI productivity follows a "J-curve", where upfront intangible investments in process redesign, business models, and human capital initially depress measured productivity before delivering substantial long-term gains³. Robert Solow's remark that "you can see the computer age everywhere but in the productivity statistics"⁴ can be expected to hold true for AI adoption as well, although the scale and timelines will differ.

Advancing frontier innovation and diffusion require different policy priorities. Frontier innovation requires concentrated, highly-skilled research talent, large risk capital, and access to advanced computing infrastructure. Diffusion depends more on widely distributed skills, organisational capabilities, interoperability, credible evidence about what works, clear rules on liability, and affordable supporting infrastructure.

Within the AI stack, frontier innovation focuses more on building more capable models and the infrastructure needed to train and run them. Diffusion requires a broader ecosystem of tools, software, standards, data and technical capabilities that allow organisations to adopt, adapt and deploy AI. Seger and Ward-Jackson from the Tony Blair Institute illustrate this with a transport analogy: middle powers can capture the benefits of AI by building the "roads" that put models to use, without having to build the most powerful "engines" themselves⁵. Since both

priorities compete for limited fiscal resources and administrative capacity, governments must make deliberate choices to balance them.

3 Why India Should Accelerate AI Diffusion^A

The case for accelerating diffusion rests on three opportunities: increasing productivity of small firms, helping public officials deliver services, and applying India's experience in IT services to AI deployment.

3.1 Productivity of Indian firms

Smaller enterprises may lack the resources to employ specialists in marketing, software development, or customer support. AI tools can help staff carry out some of these tasks. In a 2024 survey of more than 5,000 small and medium-sized enterprises across seven OECD countries, marketing and sales, and IT were the most common business support uses of generative AI. Among users that had experienced a skills gap, 39% reported that generative AI helped compensate for it. These findings are self-reported and indicate potential benefits; they do not establish comparable gains for Indian firms.⁶

The evidence shows that gains vary substantially across workers and firms. A study conducted by Brynjolfsson et al. finds a 15% average increase in issues resolved per hour among customer-support agents, with substantial variation across workers⁷. A Kenyan experiment found no statistically significant average effect on revenues and profits. Initially stronger businesses gained, while weaker businesses experienced poorer outcomes. These results suggest that cheaper access to advice does not remove the need to judge when and how to use it⁸. Bloom et al. show that better management practices alone significantly boosted productivity across large Indian textile firms. Their findings point to the value of improving everyday management, although they do not establish whether AI can deliver those improvements⁹.

Indian firms also face competitive pressure to adopt AI as it allows them to develop new services and respond to changing demand. Slower adoption could weaken their position in the global market. This matters particularly for tradable services: India's services exports reached 10.0% of GDP in H1 FY26, with software services accounting for over 40% of the total¹⁰.

The ILO's 2025 assessment notes that, as firms adopt AI, job transformation is the more likely overall outcome rather than jobs being displaced altogether¹¹. For India's tradable services, employment effects will depend on how firms reorganise work, whether productivity gains generate additional demand, and

In customer support, the largest gains went to less experienced and lower-skilled workers; in the Kenyan experiment, initially weaker firms fared worse. Who benefits may depend partly on how much judgement the task leaves to the user.

whether workers can move into changing roles. Diffusion policy should combine implementation support with monitoring impact on workers and providing opportunities for them to acquire relevant skills.

3.2 Public service delivery

Pritchett characterises India as a “flailing state” leading to a persistent disconnect between policy formulation and implementation¹². AI offers an opportunity to address some of the constraints that contribute to this gap. In delivery of social services, the World Bank finds that countries are deploying AI in front-end applications and back-end administrative applications¹³. AI can help frontline officials access relevant guidance, process documents, and communicate with citizens in their own languages. Within administrations, it can support forecasting, manage workflows, and detect irregularities in procurement or expenditure. This creates an opportunity to expand service capacity without requiring a proportionate increase in personnel.

For many departments, the priority is to digitise records, improve their accuracy and connect information across agencies. The World Bank identifies data quality and availability as the most commonly-reported barriers to government AI use¹⁴. The OECD also identifies staff capabilities, procurement, governance, and digital infrastructure as requirements for effective deployment. Common data standards and interfaces can reduce the work needed to integrate information across systems¹⁵.

India’s digital public infrastructure provides a foundation for supplying some of these complements. Reuse of modular building blocks such as identity, payments, and data exchange services already exist: Aadhaar based authentication is quite common, UPI based payments are ubiquitous, and Digilocker connects document issuers and organisations requesting verified documents. Building on such services, an AI application can retrieve or authenticate information and initiate an authorised workflow through an existing interface. Doing so would still require accurate records, appropriate access controls and integration with departmental processes.

3.3 India has substantial capabilities with which to pursue this opportunity

India’s large IT services industry has accumulated experience in integrating technologies, adapting software, and managing organisational change. These capabilities are relevant because productive AI deployment involves connecting models to operational data, redesigning workflows, testing performance, and supporting users. NASSCOM’s FY2025 estimates placed technology exports at approximately \$224 billion, against domestic revenues of roughly \$58 billion, illustrating the scale

and international orientation of this industry¹⁶. These capabilities are also backed up by Anthropic's country brief that shows that India ranks second globally in share of [Claude.ai](#) consumer usage and 1st in the share of AI use devoted to software-related tasks¹⁷.

Domestic diffusion deployments could give Indian firms experience in adapting AI to local languages, incomplete records and constrained service delivery. That experience could also help them serve other countries facing similar problems.

3.4 But, rapid AI diffusion also carries significant risks

In areas where it is unreliable, rapid AI diffusion can be ineffective or actively harmful. Dell'Acqua et al. coined the term "jagged frontier" to describe why AI excels at some tasks while failing at others in unpredictable ways¹⁸. In experiments with BCG consultants, they found that using AI for tasks in areas where it is unreliable can lead to performance worse than a human-only baseline. Additionally, Indian deployments will frequently involve models trained on non-representative data, and serving populations with low ability to contest an authoritative-sounding output. These make the risks higher.

Diffusion policy should also prioritise tangible outcomes, rather than focusing on the technology itself. Targets should be expressed in the units the sector actually cares about, and measured against a counterfactual, such as diagnostic accuracy in primary health centres, crop yield per hectare, learning outcomes, or case disposal times.

4 The Limitations Of The No Intervention Path

Diffusion will happen in India without state intervention. However, the shape and speed of the unguided path are unlikely to serve most of the population effectively.

Blume Ventures' Indus Valley reports segment the Indian population into India1, the affluent 120–140 million who constitute the addressable market for most startups; India2, an emerging aspirant class of roughly 300 million; and India3, the remaining billion or so with minimal discretionary income. Markets address the needs and wants of India1 comprehensively and India2 partially. Many people in India3 cannot pay enough to make services commercially viable, even where those services could bring them substantial benefits. The digital divide that accompanied IT proliferation could recur as a cognitive divide: unequal access to the tools of learning, judgment, and personal productivity, which compounds across a lifetime¹⁹.

The World Bank distinguishes adoption, adaptation, and

India ranks second globally in total Claude.ai use, but 101st out of 116 countries when use is adjusted for the working-age population. Both can be true: a large, technically sophisticated user base can coexist with limited penetration across society.

advancement as three ways to benefit from AI²⁰. Reaching India3 will require adaptation to deliver services in local languages, delivery through institutions people already use, and support for users who cannot pay the full cost. Commercial providers have stronger incentives to serve India1, whose users can pay.

It is not just a question of affordability. Inference prices have been falling steeply over the past few years. For instance, the cost of matching GPT-4's performance on graduate-level science questions dropped roughly 40× per year, and across benchmarks the decline ranges from 9× to 900× annually²¹. This is equivalent to a drop from \$40 to \$1 for the same volume of tokens at comparable benchmark performance over a year, although it is unclear if this trend will continue.

What India3 lacks are things that must exist around the model for it to produce value: data that represents their conditions, interfaces in their language and modality, workflow integration in the institutions they actually interact with (the PHC, the mandi, the panchayat office, the district school), standards for how a private player may connect and interact with the systems, a trusted intermediary, and recourse when the system is wrong. Providers may underinvest in these complements where users have little purchasing power or where benefits extend beyond the paying customer. Public support should target those gaps.

India's linguistic diversity presents a significant challenge, as model quality degrades sharply outside high-resource languages. The degradation is not merely lower fluency, but also an increased tendency for hallucinations and cultural meaninglessness, in ways that are hardest to detect for exactly the users least equipped to challenge the output.

Rohera et al. found that LLMs frequently outperform in English, even for questions rooted in Indic contexts, while exhibiting a greater tendency for hallucinations in low-resource Indic languages²². Increased AI adoption in Japan and Korea were associated with the release of models that performed substantially better in regional languages²³, demonstrating that language quality is an important diffusion variable. The vast diversity of India's languages presents a challenge that is unlikely to be addressed by markets as the rewards might be small compared to efforts involved. The government already addresses such gaps through efforts such as BHASHINI, which provides translation, speech recognition and speech synthesis services that other platforms can integrate²⁴. The policy priority is to build on such initiatives by funding resources for underserved languages and dialects, testing performance on actual service tasks, and supporting adoption where commercial incentives remain weak.

Voice interfaces could extend AI services to users poorly served by text-based systems, particularly in health, agriculture, and welfare²⁵. They also make data protection more complex. A single interaction may pass through speech recognition, language-model and speech-synthesis services operated by

A citizen need not own a smartphone or interact with a chatbot to benefit from AI. A nurse, agricultural worker or public official could use it in ways that benefit them. This suggests measuring diffusion through improvements in the services people receive, alongside direct adoption. The intermediary's ability to interpret, challenge and act on outputs becomes crucial.

different vendors, creating several points at which recordings and transcripts are processed or retained. The Digital Personal Data Protection Act, 2023 establishes responsibilities for processing personal data, including processing undertaken on an organisation's behalf. The rules and commencement notifications issued in November 2025 provide for phased implementation of this framework^{26, 27}.

Applying these requirements to voice services calls for practical guidance. Beyond information about a person's context and needs, voice is a biometric marker that might convey information about their health and emotional state²⁸. Providers need usable ways to explain data practices in local languages, obtain and record consent where it is the basis for processing, and support withdrawal and deletion across the service chain, subject to applicable retention requirements.

Guidance should distinguish processing needed to deliver a service from subsequent use of recordings or transcripts for model training, and clarify responsibilities where vendors determine their own purposes for using the data. Addressing these implementation questions would reduce uncertainty for providers, and help users exercise meaningful control over their information.

5 The Four Conditions for AI Diffusion

This paper proposes four conditions for beneficial and sustained AI diffusion: reliability at the task, access to the service, complementary capital, and permissions. They can be used to assess a particular deployment and identify what is preventing it from delivering benefits.

The examples below apply the four conditions to specific deployments: human-reviewed court transcription and translation; voice-based crop advice for a defined crop, region and season; and clinician-facing diagnostic support for adults presenting with acute fever at a primary health centre. Each example specifies an intended outcome and questions for assessing the constraints.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
AI produces drafts that authorised staff verify before use. The objective is to reduce the time and cost of producing accurate, accepted records and translations. Transcription and translation require separate evaluation.	A telephone service provides locally relevant crop-management advice with specialist escalation. The objective is timely adoption of appropriate practices and improved net farm returns.	AI supports a clinician's assessment and referral decisions. The objective is better clinical management, accounting for missed serious illness, unnecessary treatment, and staff time.

5.1 Reliability at the task

Does the deployed system meet an acceptable performance standard under actual operating conditions? Reliability concerns the complete system, including its data sources, human review and fallback arrangements. Strong average model performance may conceal significant failures for particular tasks, languages, or populations.

Bottlenecks

- Performance gaps across local populations, dialects, environments, and operating conditions.
- Missing, inaccurate, or outdated contextual information.
- Recurrent errors obscured by aggregate statistics.
- Unclear or porous process and decision boundaries that lead to:
- Failure to recognise uncertainty, remain within a suitable task boundary, or escalate appropriately.
- Verification burdens that erase productivity gains, and performance deterioration following changes to models or their operating environment.

Policy instruments

- Independent, task-specific evaluation against existing practice, with results disaggregated by relevant population and operating conditions.
- Shared evaluation datasets, representative data collection, and maintained open domain knowledge sources.
- Procurement requirements for uncertainty handling, human escalation and fallback.
- Post-deployment monitoring, incident reporting, and revalidation after material changes.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
Does the system preserve names, numbers, negation, speaker identity and legal meaning across courtroom conditions and language pairs? Which consequential errors persist after review?	Is advice appropriate to crop stage, soil, weather and available inputs? Distinguish errors in recognising speech from errors in crop recommendations.	Does the combined clinician–AI workflow perform adequately on local patients using measurements actually available at the facility? Assess missed serious cases and false alarms.

5.2 Access to the service

Can intended users obtain and use the service affordably, when and where they need it? Access ranges from the infrastructure supplying computation to the channels through which people use the service.

Bottlenecks

- Service prices and recurring delivery costs.
- Constraints on energy, compute, devices, and connectivity.
- Excessive latency or downtime, or insufficient capacity.
- Interfaces that exclude users because of language, literacy, disability or lack of assistance.
- Dependence on a single supplier or platform that threatens affordability or continuity.

Policy instruments

- Pooled procurement and targeted service subsidies where affordability limits beneficial use.
- Investment in power resilience, connectivity, and shared inference infrastructure where these constrain delivery.
- Telephone, low-bandwidth, offline, and assisted channels suited to the intended population.
- Accessible, multilingual interfaces tested with intended users.
- Portability requirements and competitive procurement to reduce switching barriers.

Availability of a model or application alone does not establish effective access. Open-weight and smaller models are possible delivery strategies, rather than universal solutions. Their suitability depends on task performance, maintenance requirements, and total operating costs.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
Can targeted courts capture usable audio and obtain the service consistently? Can intended readers access the resulting records?	Who can complete a call at an acceptable cost and at the relevant time, including users of shared phones and unreliable networks?	Are devices, power, and connectivity available during consultations? Are response times acceptable and interruptions manageable?

5.3 Complementary capital

Can adopters and their delivery networks integrate the system, act on its outputs, and sustain the gains? Complementary capital includes operational data, skills, organisational routines, incentives, and the resources needed to translate an AI output into a useful outcome. These capabilities extend beyond the institution purchasing the system. Farmers may need inputs and finance to act on advice; clinicians may need functioning testing, treatment and referral services.

Training staff, cleaning records, and redesigning workflows consume resources before they produce measurable gains. The productivity J-curve describes how these investments can initially depress measured productivity, with benefits appearing later. Early disappointment therefore calls for diagnosis: an organisation may be building useful capabilities, or it may be funding an ineffective deployment. These reforms can also improve productivity without AI. Where feasible, evaluations should compare three groups: current practice, the same data and workflow improvements without AI, and those improvements with AI. Where such comparisons are infeasible, findings should be described as the effect of the combined intervention, without attributing the entire gain to AI.

A low baseline of productivity does not automatically guarantee large feasible gains from AI. AI cannot resolve a shortage of medicines, irrigation or working capital. The relevant comparison is the best feasible alternative use of existing resources, including simpler software, staff training, process reform or improved physical services.

Bottlenecks

- Fragmented records, poor interoperability and costly integration with existing systems.
- Insufficient staff capability, implementation time, and workflow ownership.
- Weak procurement, maintenance, and vendor-management capacity.

- Misaligned incentives, including situations where the adopter bears costs but benefits accrue elsewhere.
- Inadequate transition funding and recurring budgets.
- Missing downstream resources needed to act on recommendations.

Policy instruments

- Interoperable standards, usable data interfaces, and shared implementation services, including appropriate connections to digital public infrastructure.
- Deployment teams, task-specific training, and protected time for learning and workflow redesign.
- Stronger technology leadership and procurement capacity in public institutions.
- Transition funding, sustainable operating budgets, and revised responsibilities or incentives.
- Investment in the extension, referral, input-supply and other services necessary to realise the intended benefit.

Purchasing software and conducting introductory training will not, by themselves, establish the working practices and supporting services needed to use AI effectively. Adopting institutions should assign responsibility for implementation, give staff time to test and revise workflows, and budget for maintenance and continuing support. They should also check whether the people receiving AI-assisted advice can act on it. Funding and evaluation should therefore cover the full delivery process, including the services and resources needed to turn recommendations into outcomes.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
Can judges, staff, lawyers and litigants use permitted records to prepare orders or submissions faster? Are review, filing and version control integrated? Does it deliver gains after review and verification costs are considered?	Can the provider maintain local information and handle escalations? Can farmers obtain the labour, finance, water and inputs needed to act?	Does the system fit consultations and recordkeeping? Can staff complete recommended tests, treatment and referrals?

5.4 Permissions

Is the deployment legally authorised, institutionally accepted, and sufficiently trustworthy for appropriate use? Permissions include both regulatory compliance and the accountability arrangements that build trust in the system.

Bottlenecks

- Unclear permitted uses, professional boundaries, and decision rights.
- Unresolved responsibilities among developers, service providers, and adopting institutions.
- Uncertainty about lawful data use, privacy, and confidentiality.
- Inadequate procedures for challenging outputs, correcting errors, and obtaining redress.
- Distrust arising from opacity, exclusion, conflicting commercial incentives, or previous failures.

Policy instruments

- Sector-specific guidance and clear approval pathways for defined uses.
- Explicit allocation of responsibilities, supported by audit trails and incident procedures.
- Supervised sandboxes with evaluation requirements and clear conditions for progression or exit.
- Understandable disclosure and accessible correction, contestation and redress.
- Participation by frontline workers and affected users in system design, testing and decisions about deployment.
- Stronger regulatory and institutional capacity to assess deployments and oversee their operation.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
Who authorises recording, certifies outputs, and resolves disputed text? What confidentiality rules and restrictions on use apply?	Are data practices understandable and accountable? Are the adviser's commercial incentives transparent, and is responsibility for harmful advice clear?	Is the intended use authorised? Are clinician and supplier responsibilities, patient-data practices, and incident procedures clear?

The four conditions describe requirements for effective diffusion, although systems can spread despite shortfalls in some areas. Widespread adoption alone is therefore not evidence that the conditions have been met.

The conditions are also interdependent and involve tradeoffs in some cases. Human review may improve reliability while increasing demands on staff and budgets. A voice interface may improve access while introducing speech-recognition errors. Clear accountability may strengthen institutional acceptance and require incentives to monitor performance.

The examples show why improving access alone may produce little benefit. A court may have a capable transcription service but lack time to verify its output; a clinician may have diagnostic support but no way to arrange the recommended tests. Diffusion policy must address the constraint that limits each deployment.

Court transcription and translation	Crop advisory over voice	Diagnostic support at a PHC
Local performance and review capacity may limit gains. Authorisation depends on the intended use of the output.	Local reliability and the capacity to act may limit benefits. Access and trust may constrain particular groups.	Clinical validity and workflow capacity may interact with limitations in testing, treatment and referral. Access and authorisation require separate verification.

5.5 Mapping the conditions to the AI stack

Jensen Huang characterises AI infrastructure as a five-layer cake spanning energy, chips, infrastructure, models, and applications²⁹. In addition, data can be thought of as a cross-cutting input for building and evaluating AI models and applications. Mapping the four diffusion conditions against this stack helps identify what is required at each layer to deliver the services.

Diffusion condition	Relationship to the layers of the AI stack	Procurable inputs and capabilities that must be developed
Reliability at the task	Primarily data, models and applications, supported by dependable infrastructure.	Models, applications, evaluation services, and monitoring tools can be procured. Representative data, local validation, and responsibility for performance must be established.

Diffusion condition	Relationship to the layers of the AI stack	Procurable inputs and capabilities that must be developed
Access	All layers. Extends to devices, connectivity and delivery channels.	Capacity and services can be purchased. Affordability, accessibility, and competition require design and governance.
Complementary capital	Primarily resides in organisations and delivery networks beyond the technical stack. Within the stack, relates especially to data and applications, and their integration with existing systems.	Integration, training, and implementation support can be purchased. Effective processes, incentives, and operational knowledge must be developed.
Permissions	Cross-cutting across infrastructure approvals, model and data rights, and sectoral authorisation and accountability for applications.	Compliance and assurance services can be procured. Legal authority, governance, and trust cannot be obtained through procurement alone.

Sovereignty debates often focus on dependence on foreign chips, computing infrastructure and models. These dependencies matter for access. Reducing them would still leave the problems of local reliability, organisational capability, and accountability unresolved, which are essential for diffusion.

6 What India Should Do To Drive AI Diffusion?

India's diffusion strategy should build on the programmes and institutions already established or envisaged under the IndiaAI Mission. It includes support for application development, skills, data access, and safe and trusted AI alongside computing infrastructure and domestic model development³⁰. The India AI Governance Guidelines, published in 2025, advocate empowering sectoral regulators and envisage technical support from the AI

Safety Institute³¹. The recommendations below would give these efforts a more explicit focus on the conditions for beneficial diffusion.

6.1 Reliability at Task

6.1.1 Establish common evaluation standards and deployment-specific validation

The government should use public procurement requirements and sectoral standards to promote consistent evaluation methods and standardised reporting on AI performance, costs and failure modes. Procurement rules should require suppliers to provide evaluation evidence proportionate to the intended use and risk of the system, using common reporting formats that allow findings to be compared and reused where relevant. This would build on the India AI Governance Guidelines, which envisage the AI Safety Institute developing evaluation metrics and testing frameworks in collaboration with relevant agencies and sectoral regulators³².

To extend this work, the government should establish an accreditation mechanism for independent AI evaluators. The national AI Safety Institute and sectoral regulators should coordinate the technical criteria for accreditation and how it will be conducted. Accredited academic and commercial laboratories could then conduct evaluations within specified areas of expertise. Evaluators should be reassessed periodically, and required to disclose and manage conflicts of interest.

Accreditation should establish an evaluator's competence for a defined scope of testing. Suppliers and adopting institutions should remain responsible for validating each deployment, including its data sources, human review, escalation and fallback arrangements. Public support should prioritise evaluations with substantial learning spillovers or benefits for underserved users, with publication of reusable findings where confidentiality permits.

India should also work with AI Safety Institutes in other countries to build consensus on evaluation methods and deployment-specific validation requirements. Alignment could support mutual recognition of evaluation results where standards and testing scopes are comparable. It could also create opportunities for India's IT services industry to provide evaluation and validation services globally, drawing on its experience in integrating and adapting software systems to different organisational and sectoral contexts.

6.1.2 Require evidence of performance across regional languages and user groups

Public procurement requirements and sectoral evaluation standards should require AI services to demonstrate performance in the languages, dialects, and user groups they are intended to serve. Evaluations should reflect actual tasks and conditions of

use, with results reported separately for relevant groups so that aggregate scores do not conceal gaps.

A consortium of IITs, IIITs, university language departments, and regional academic institutions could help develop evaluation methods and test sets, drawing on linguistic expertise across the country. Participation should reflect the languages and dialects being assessed and include speakers from the intended user communities.

Datasets commissioned through public procurement should include usable documentation and clear reuse terms, with controlled access where privacy or other rights prevent open release. Success should be measured by improved task completion and fewer errors for intended users, alongside reuse of evaluation methods and datasets across services.

6.1.3 Tie procurement and expansion to evidence of useful outcomes

Contracts should define acceptable performance before deployment and fund testing against current practice and the best feasible alternatives. Net benefits should be measured after verification, correction, training, and operating costs, as gains in some areas might be balanced by the need for human verification.

6.1.4 Treat reliability as a continuing obligation

Adopting institutions should assign responsibility for system performance, validate systems under conditions representative of their intended use, and monitor outcomes after deployment. NIST's AI Risk Management Framework provides a useful reference for these practices, including documented responsibilities, evaluation in deployment conditions, and procedures for monitoring, incident response, and recovery³³.

Public procurement contracts should translate these practices into specific obligations. They should define reportable incidents, response times, conditions for suspension, and workable fallback arrangements. Suppliers should disclose material updates, and adopting institutions should require revalidation when changes to models, data, task scope, or operating conditions could affect performance. Appropriately anonymised findings from unsuccessful deployments should be published alongside successes, so that other adopters can learn from both.

6.2 Access to the Service

6.2.1 Make open-weight models a default option in public procurement and ensure interoperability and portability

Open weight models reduce dependence on a model provider's continued API access and commercial terms. This can strengthen

sovereignty without requiring every model to be developed domestically³⁴. Strengthening the open AI ecosystem will also have spillover effects on the wider community, and public procurement could be a means to achieve that. To use a proprietary model should require the provider to demonstrate that the model is significantly better for a given task than comparable open models within time and cost constraints.

Procurement contracts should also preserve the ability to change models and suppliers through interoperable interfaces, exportable data and configurations, and clear migration rights. These requirements should apply regardless of whether the selected model is open or proprietary and it should not require rebuilding the service. Memory and institutional context should be retained with the institutions deploying the systems and they should be able to transfer it and migrate to a different vendor if required. In other words, changing vendors should not mean losing valuable institutional context.

6.2.2 Support an international consortium to sustain competitive open models

India should help establish or join a consortium of companies, research institutions and partner countries that pools funding, compute and expertise to develop and maintain open model families. Nathan Lambert argues that collective industry funding could provide a more durable foundation for near-frontier open models as individual firms' incentives to release them change³⁵. Switzerland's Apertus initiative illustrates collaboration between public research institutions, shared computing infrastructure and industry partners³⁶. The French company Mistral intends to make sovereign open-weight AI that close the gap with frontier³⁷. NVIDIA's Nemotron Coalition illustrates an industry initiative among companies and model developers³⁸. There is an alignment of interests between multiple countries and companies in building a capable open AI ecosystem.

India should participate by providing funding and expertise in evaluating, adapting, and deploying AI systems. The commitments should support successive model releases, ongoing maintenance, and adaptation to different languages and sectoral needs. The consortium should require participating developers to release model weights, documentation, and supporting tools under licences that allow broad use, modification, and redistribution. This would spread development and maintenance costs across participating countries and companies to ensure a dependable supply of capable models they can use and adapt.

6.2.3 Pool public-sector inference demand while preserving choice and accountability

National-scale diffusion will require reliable inference capacity and a practical way for institutions to purchase it. A shared

public-sector platform, similar to OpenRouter, could let departments buy inference from multiple approved providers, with common arrangements for billing, logging, evaluation, and data handling. This could reduce repeated procurement, security, and evaluation work across departments. Common interfaces would make switching providers easier, although a replacement model would still need to be tested for the intended use. Vendors could be empanelled against published requirements, audited periodically, and removed if they fail to meet those requirements.

6.2.4 Reduce the friction in buying technology from abroad

As Raman and Shah argue, India should review the barriers to purchasing IT equipment and overseas technology services so that firms and individuals can readily participate in the global technology ecosystem³⁹. This means simplifying cross-border payments, enabling routine credit-card purchases of cloud and API services, and reducing avoidable customs delays and documentation burdens. Clear compliance requirements and effective resolution of recurring payment failures would particularly benefit small firms and independent developers. For a small firm or an individual developer, a blocked card or a customs delay is a harder barrier than the price of tokens.

6.2.5 Make underserved AI applications affordable through existing service networks

The government should enable farmer producer organisations, cooperatives, and other community institutions to arrange shared access to independently-evaluated AI services. For example, a farmer producer organisation could subscribe to crop advisory services in its members' languages, and deliver crop advice to farmers through existing field staff and communication channels. Collective subscriptions could spread costs across users and give these organisations greater bargaining power.

Where members cannot afford services with demonstrated benefits, existing sectoral support budgets or philanthropic partnerships could help cover subscription and delivery costs. Participating organisations should retain the ability to select and switch providers, subject to transparent selection and published quality requirements. Oversight should focus on independently assessed service quality, evidence that intended users are being reached, and periodic assessment of benefits.

6.2.6 Treat electricity distribution reform as AI strategy

Reliable, competitively priced electricity is essential for AI infrastructure at scale. State-sector DISCOMs had accumulated losses of ₹7.08 lakh crore in FY2023–24, underscoring the need to improve their financial and operational performance alongside investment in power supply for data centres⁴⁰.

Near-term action should build on the Electricity (Amendment) Rules, 2026, which ease captive-power arrangements through clearer ownership provisions, more flexible group captive structures and streamlined verification. States should translate these changes into timely grid connections, predictable charges and consistent administration, allowing AI infrastructure developers to use captive generation more effectively. Broader reforms proposed in the Draft Electricity Amendment Bill, 2025 remain a legislative agenda⁴¹.

Andhra Pradesh's April 2026 policy proposing deemed distribution licences for qualifying data centres is a contested example of efforts to enable dedicated power arrangements⁴². Its legal basis has been challenged: the Human Rights Forum argues that self-consumption does not qualify as distribution and that an executive order cannot create a new class of deemed licensee outside the statutory framework⁴³. Dedicated supply arrangements are worth considering, but any licensing route should have a clear legal basis and required regulatory approvals.

6.3 Complementary capital

6.3.1 Build practical implementation capabilities and budget for organisational change

Industry associations, universities, and private service providers can help smaller firms identify useful AI applications, prepare data, redesign workflows and assess results. Shared advisory services and peer learning can make this expertise affordable without each firm hiring a specialist team. This approach reflects Mazarr's emphasis on widely distributed implementation skills and learning among practitioners⁴⁴ and Brynjolfsson et al.'s account of the intangible investments needed to realise productivity gains from general-purpose technologies⁴⁵.

Adopting organisations should budget for integration, staff learning, maintenance, and evaluation alongside software and compute. These investments take time, and productivity gains may emerge only after working practices change. Government support should be limited to selected pilots whose findings can benefit other adopters, with publication of both positive and negative results.

6.3.2 Invest in the skills and judgement needed for AI adoption

India's diffusion strategy should strengthen the broad base of people who can adopt, adapt and advance AI systems^{46, 47}. This builds on the role that institutions such as IITs, IISc, NCST, and ERNET played in developing capabilities for India's IT industry⁴⁸.

Universities, vocational institutions, and employers should use tasks people encounter at work to teach them to check AI outputs, recognise errors, and decide when to seek specialist

help. Industry-academia partnerships and mid-career courses can connect learning to workplace needs. Contracting with universities and private firms offers a way to effectively build capabilities across the economy, as proposed by Mashelkar, Shah and Thomas⁴⁹.

6.3.3 Invest in open, common standards to enable vibrant AI markets

Open standards and interfaces to relevant digital public infrastructure can help systems exchange information reliably. They allow different providers to build compatible services over the standard interfaces allowing customers to switch suppliers. India's experience with identity (Aadhaar) and payments (UPI) systems provides a strong foundation to shape global standards for how AI systems can interact with these types of systems.

Organisations must also be able to act on AI outputs: farmers need access to inputs and finance to follow crop advice, while clinicians need testing, treatment and referral services to use diagnostic recommendations. DPI's exclusion, privacy, and surveillance risks scale with AI integration rather than remaining constant, and the consent architecture and safeguards must be strengthened in parallel.

6.4 Permissions

6.4.1 Build institutional trust and give officials room to learn

AI adoption, especially in public service delivery, should involve frontline workers and affected users in shaping how systems are introduced and used. This can help resolve the conflicts with existing practices or challenges in access and grievance redressal, and drive acceptance among users. Public institutions should also give officials the authority to test promising applications within clear limits, recognising that a carefully conducted experiment that finds no benefit still produces valuable knowledge. Testing responses to practical problems and adapting them in light of results aligns with the "Problem-Driven Iterative Adaptation" approach to building capacity⁵⁰.

6.4.2 Make data-protection requirements easier to apply in practice for AI use cases such as voice

Industry bodies and service providers should develop practical examples of applying the DPDP framework to AI services, with regulatory clarification where recurring uncertainty remains. Priorities include understandable notices and usable consent and withdrawal processes for voice-based services, including in local languages and dialects. Examples should explain how to distinguish data needed to deliver a service from subsequent use for training or personalisation, and how responsibilities apply

when several providers handle the same data. This support should help organisations comply as the framework takes effect.

6.4.3 Strengthen sectoral regulation with dedicated expertise and funding

India should build on existing laws and sectoral regulators rather than introduce a single omnibus AI Act. AI's risks depend heavily on its context of use, making sectoral institutions well placed to assess acceptable performance, professional responsibilities and potential harms. Their effectiveness, however, depends on having the capacity to exercise these responsibilities. Expanding regulatory mandates should therefore be accompanied by dedicated, predictable budgets for technical recruitment, staff training, independent evaluation, and enforcement. Shared technical expertise can support regulators facing common challenges, while sustained funding helps them retain institutional knowledge and assess industry claims independently. A coordinated national framework should establish common principles and clarify responsibilities.

7 Conclusion

India should accelerate AI diffusion to raise firm productivity and improve public services. Achieving these gains requires more than access to capable models. Systems must work with local data, available staff, and the institutions responsible for delivering services. The four conditions developed in this paper help identify where a deployment falls short and the proposed policy interventions address constraints that adopters and commercial providers are unlikely to resolve on their own.

Widespread AI adoption may nevertheless come with job losses and reduced demand for some skills. Where these costs are concentrated, they can cause substantial hardship and generate political opposition to adoption. Employers will need to help staff move into changing roles, and governments will need to support those who lose out. Although this paper has focused on making AI deployments effective, sustaining diffusion will also require supporting the workers who bear the costs of transition.

Endnotes

1. Mazarr, Michael J. A New Age of Nations: Power and Advantage in the AI Era. Expert Insights PE-A3691-14. RAND Corporation, 2026. [Link](#)
2. Paul A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," American Economic Review 80, no. 2 (May 1990): 355–61, [Link](#).
3. Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics 13, no. 1 (2021): 333–72, [Link](#).
4. Jack E. Triplett, "The Solow Productivity Paradox: What Do Computers Do to Productivity?," Brookings Institution, March 1, 1999, [Link](#).
5. Tony Blair Institute for Global Change, "Build the Roads, Not the Engine: An Open-Source Strategy for Middle-Power AI Sovereignty," September 14, 2026, [Link](#).
6. Marguerita Lane and Carla Ruggiu, Generative AI and the SME Workforce: New Survey Evidence (OECD Publishing, 2025), [Link](#).
7. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work," Quarterly Journal of Economics 140, no. 2 (2025): 889–942, [Link](#).
8. Nicholas G. Otis et al., "The Uneven Impact of Generative AI on Entrepreneurial Performance: Evidence from a Field Experiment in Kenya," Management Science (forthcoming), pre-published July 10, 2026. [Link](#).
9. Nicholas Bloom et al., "Does Management Matter? Evidence from India," Quarterly Journal of Economics 128, no. 1 (2013): 1–51, [Link](#).
10. Ministry of Finance, Department of Economic Affairs, "Services: From Stability to New Frontiers," in Economic Survey 2025-26 (Government of India, 2026), 259, [Link](#).
11. Pawel Gmyrek et al., "Generative AI and Jobs: A Refined Global Index of Occupational Exposure," ILO Working Paper 140 (International Labour Organization, 2025), [Link](#).
12. Lant Pritchett, "Is India a Flailing State? Detours on the Four Lane Highway to Modernization," HKS Faculty Research Working Paper RWP09-013, May 2009, [Link](#).
13. World Bank, World Development Report 2026: The Promise of Artificial Intelligence (World Bank, 2026), 156–57. [Link](#).
14. World Bank, World Development Report 2026: The Promise of Artificial Intelligence (World Bank, 2026), 282–86. [Link](#).
15.
 5. OECD, "Enablers, Guardrails and Engagement for Unlocking Trustworthy AI," in Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions (OECD Publishing, 2025), 116–21, [Link](#).
16. NASSCOM, "Strategic Review Press Release" (National Association of Software and Service Companies, 2025), [Link](#).
17. Ruth Appel, "India Country Brief: The Anthropic Economic Index," Anthropic, February 16, 2026, [Link](#).
18. Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, et al. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality." Harvard Business School Working Paper 24-013, September 2023. [Link](#).
19. Pai, Sajith, Anurag Pagaria, Nachammai Savithiri, and Dhruv Trehan. Indus Valley Annual Report 2025. Blume Ventures,

2025. [Link](#).

20. World Bank, World Development Report 2026: The Promise of Artificial Intelligence (World Bank, 2026), 233–234. [Link](#).
21. Ben Cottier, Ben Snodin, David Owen, and Tom Adamczewski, “LLM Inference Prices Have Fallen Rapidly but Unequally across Tasks,” Epoch AI, March 12, 2025, [Link](#).
22. Pritika Rohera et al., “Better to Ask in English? Evaluating Factual Accuracy of Multilingual LLMs in English and Low-Resource Languages,” arXiv, September 14, 2025, [Link](#).
23. Microsoft Research. Global AI Diffusion: Q1 2026 Trends and Insights. Microsoft Corporation, 2026. [Link](#).
24. PIB Research Unit, “Artificial Intelligence for Culture and Languages,” Press Information Bureau, Government of India, February 9, 2026, [Link](#).
25. Shweta Mohandas et al., Making Voices Heard (Centre for Internet and Society, 2021), [Link](#).
26. The Digital Personal Data Protection Act, 2023, No. 22 of 2023, Ministry of Electronics and Information Technology, [Link](#).
27. Ministry of Electronics and Information Technology, Digital Personal Data Protection Rules, 2025, , Government of India, [Link](#).
28. Marie-Françoise Malo et al., “Using Voice and Speech Data in Healthcare: A Scoping Review of the Ethical, Legal and Social Implications,” Frontiers in Digital Health 8 (2026): 1750111, [Link](#).
29. Jensen Huang, “AI Is a 5-Layer Cake,” NVIDIA Blog, March 10, 2026, [Link](#).
30. Prime Minister’s Office, “Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem,” Government of India, March 7, 2024, [Link](#).
31. India AI Governance Guidelines Committee, India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation (Ministry of Electronics and Information Technology, November 2025), [Link](#).
32. *ibid*.
33. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (U.S. Department of Commerce, January 2023), [Link](#).
34. Pranay Kotasthane, Nitin Pai, and Bharath Reddy, “Why Is China Giving Away Its AI Models? Five Possible Explanations for China’s Open-Weight Strategy,” Takshashila Discussion Document 2026-25 Takshashila Institution, 2026, [Link](#).
35. Nathan Lambert, “The Inevitable Need for an Open Model Consortium,” Interconnects, April 11, 2026, [Link](#).
36. CSCS, ETH Zurich, and EPFL, “Apertus: A Fully Open, Transparent, Multilingual Language Model,” Swiss National Supercomputing Centre, September 2, 2025, [Link](#).
37. Mistral AI, “Mistral Raises €3B to Make Sovereign, Open-Weight AI the Technology Frontier,” September 8, 2026, [Link](#).
38. NVIDIA, “NVIDIA Launches Nemotron Coalition of Leading Global AI Labs to Advance Open Frontier Models,” NVIDIA Newsroom, March 16, 2026, [Link](#).
39. Siddharth Raman and Ajay Shah, “An Indian Strategy for AI,” Business Standard, June 22, 2026, [Link](#).
40. Sixteenth Finance Commission, Report for 2026–31, Volume I: Main Report (Government of India, 2025), chap. 13, [Link](#).
41. Press Information Bureau, “Government Amends Electricity Rules to Strengthen Captive Power Framework and Support Industrial Competitiveness,” Ministry of Power, Government of India, [Link](#).

42. "What Google's Discom Licence Tells Us about the Future of Energy in India," ETEnergyworld, June 13, 2026, [Link](#).
43. Human Rights Forum, "AP Govt's Data Center Deemed Distribution Licence Policy (DDL) Is Legally Unsustainable, Must Be Withdrawn," June 22, 2026, [Link](#).
44. Michael J. Mazarr, "A New Age of Nations: Power and Advantage in the AI Era", Expert Insights PE-A3691-14, RAND Corporation, 2026, [Link](#).
45. Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics 13, no. 1 (2021): 333-72, [Link](#).
46. Kailash Nadh, "DeepSeek, AI Sovereignty, and India," nadh.in, January 29, 2025, [Link](#).
47. Siddharth Raman and Ajay Shah, "An Indian Strategy for AI," Business Standard, June 22, 2026, [Link](#).
48. IITs: Indian Institutes of Technology; IISc: Indian Institute of Science; NCST: National Centre for Software Technology; ERNET: Education and Research Network. ERNET began in 1986 with NCST, IISc, five IITs and the Department of Electronics as participating institutions.
49. R. A. Mashelkar, Ajay Shah, and Susan Thomas, "Rethinking Innovation Policy in India: Amplifying Spillovers through Contracting-Out," Accounting Theory and Practice 1 (December 2024), [Link](#).
50. Matt Andrews, Lant Pritchett, and Michael Woolcock, "Escaping Capability Traps through Problem-Driven Iterative Adaptation (PDIA)," CGD Working Paper 299 (Center for Global Development, June 2012), [Link](#).



TAKSHASHILA
INSTITUTION

The Takshashila Institution is an independent centre for research and education in public policy. It is a non-partisan, non-profit organisation that advocates the values of freedom, openness, tolerance, pluralism, and responsible citizenship. It seeks to transform India through better public policies, bridging the governance gap by developing better public servants, civil society leaders, professionals, and informed citizens.

Takshashila creates change by connecting good people, to good ideas and good networks. It produces independent policy research in a number of areas of governance, it grooms civic leaders through its online education programmes and engages in public discourse through its publications and digital media.

©The Takshashila Institution, 2025