



For People, Planet, and Progress: Perspectives from India's AI Impact Summit

Nidhi Singh | Tejas Bharadwaj | Shruti Mittal
Adarsh Ranjan | Konark Bhandari | Charukeshi Bhatt

For People, Planet, and Progress: Perspectives from India's AI Impact Summit

Nidhi Singh | Tejas Bharadwaj | Shruti Mittal
Adarsh Ranjan | Konark Bhandari | Charukeshi Bhatt

© 2026 Carnegie Endowment for International Peace. All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are those of the author(s) and do not necessarily reflect the views of Carnegie, its staff, or its trustees.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from Carnegie India or the Carnegie Endowment for International Peace. Please direct inquiries to:

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, D.C. 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

Carnegie India
Unit C-4, C-5 & C-6, Edenpark,
Shaheed Jeet Singh Marg
New Delhi – 110016, India
P: +011 4008687
CarnegieIndia.org

This publication can be downloaded at no cost at CarnegieIndia.org

Photo credit: Ministry of External Affairs India / Flickr / CC BY-NC-ND 4.0

Contents

00	Introduction Nidhi Singh	1
01	From Safety to Impact: Stocktaking the AI Summit Process Nidhi Singh	3
02	The AI Safety Journey From Bletchley to Delhi Tejas Bharadwaj	6
03	Understanding the Global AI Impact Commons Shruti Mittal and Nidhi Singh	11
04	The Resource Cost of India's Data Center Expansion Adarsh Ranjan	14
05	India's AI Compute Question: Strategic Leverage Over Sovereignty Moonshots? Konark Bhandari	18
06	Rethinking Inclusion in AI Charukeshi Bhatt	21
	About the Authors	24
	Notes	25
	Carnegie India	31



Introduction

Nidhi Singh

Over the past two and a half years, a series of global AI summits has traced a revealing arc. What began at Bletchley Park in 2023 as a conversation on the existential risks of frontier AI systems progressively broadened through Seoul and Paris. It arrived in New Delhi in February 2026 as something quite different: a multilateral exercise that places the real-world impact of AI at the center of the global AI governance conversation.

The India AI Impact Summit, the first in this series to be hosted by a Global South nation, was framed around the three *sutras*—or themes—of People, Planet, and Progress and seven *chakras*—or working groups—of human capital, inclusion, trusted AI, resilience, science, democratizing AI resources, and social good. In doing so, it prioritized a set of questions that the earlier, safety-oriented iterations had not.¹ How can AI reach the last mile? Who controls the foundational resources on which AI development depends? And how can governance frameworks be designed to serve the needs of developing nations, rather than being shaped exclusively by the leading developers of frontier models? These are not merely abstract questions. They have material consequences for the trajectory of AI adoption across much of the world.

The gap between those who shape AI systems and those expected to adopt them plays out across multiple domains—in the governance frameworks that determine priorities, the concentration of compute, the environmental costs of infrastructure expansion, and the linguistic foundations that determine whether AI tools work for populations beyond the English-speaking world. Understanding this divide requires sustained analytical engagement that goes beyond what any summit declaration can offer. It is this recognition that motivates the present compendium.

This collection of essays by scholars from Carnegie India's Technology and Society program opens with a stocktaking of the summit process itself. It traces the evolution of the AI summit series and examines India's framing around the three *sutras* of people, planet, and progress. Scholars have catalogued and assessed the concrete deliverables that emerged from the seven *chakras* mentioned above and assessed what the precedent of a Global South country hosting means for the future of the multilateral conversation.

The next piece examines the evolution of AI safety across the same summit arc, tracking how the discourse shifted from existential risk at Bletchley to “secure and trusted AI” in New Delhi. It also considers what this reframing means as the process moves to Switzerland.

A subsequent contribution examines one of the Summit’s flagship outcomes: the Global AI Impact Commons, a platform designed to enable the adoption, replication, and scaling of successful AI use cases across regions. We assess whether the Commons can function as a meaningful instrument for cross-border learning on AI deployment, particularly across the Global South, or whether it risks joining the growing catalogue of well-intentioned but underutilized multilateral initiatives.

The next essay turns to the material infrastructure underlying AI development, examining the wave of data center investment commitments that followed the summit. It confronts the tension between India’s ambitions for compute sovereignty and the water, energy, and grid constraints that large-scale data center expansion will impose, arguing that sustainability conditions must be embedded in policy now rather than retrofitted later.

The compendium also includes a contribution on AI hardware and the geopolitics of compute. It argues that India’s path to securing its position in the global compute hierarchy lies not in a sovereignty moonshot but in strategic leverage: identifying supply-chain bottlenecks, building industrial capacity in areas like CPU packaging and testing, and converting domestic capability into an international negotiating position.

Finally, an essay on linguistic inclusion challenges the dominant framing of AI inclusivity as a dataset problem. It argues that meaningful inclusion requires moving upstream—from who is represented in training data to who produces the language, who digitizes it, and who participates in data creation. It calls for an increased investment in the community and linguistic infrastructure that makes good datasets possible.

Together, these contributions seek to move beyond the summit communique and into the substance of equitable AI development. The India AI Impact Summit opened important avenues of conversation around the three *sutras*. In doing so, it established that last-mile diffusion and development outcomes are legitimate organizing principles for global AI cooperation, not secondary concerns to be addressed after the questions of safety.

01

From Safety to Impact: Stocktaking the AI Summit Process

Nidhi Singh

The global AI summit process has evolved considerably since its inception. Bletchley Park in 2023 focused on the existential risks of frontier AI systems, producing safety institutes and a scientific reporting mechanism. Seoul in 2024 deepened that institutional scaffolding through networked safety institutes and voluntary corporate commitments. Paris in 2025 broadened the agenda to encompass public interest, sustainability, and access to infrastructure. By the time the process reached New Delhi, the organizing frame had shifted decisively toward diffusion, adoption, and development outcomes.

The India AI Impact Summit was the first in this series hosted by a Global South nation. The scale reflected ambition: approximately 600,000 in-person attendees, delegations from more than 100 countries, twenty-two heads of state or government, and over 300 exhibitors across ten thematic pavilions. The summit was also a key moment for India to articulate its stewardship of a “third way” in AI governance, one that negotiates between the competing visions of the major AI powers and advances the priorities of the developing world.²

From Building Models to Deploying Impact

The central narrative of the Impact Summit represented a departure from earlier iterations. Previous summits engaged primarily with how to develop and govern frontier AI models safely. India sought to move the conversation toward on-ground implementation of AI. The focus was not on who is building the most capable large language model, but on how AI capabilities, once built, can be deployed at scale to address the needs of populations that existing technology paradigms have not adequately served.

This was operationalized through the summit’s distinctive architecture: one *mantra*, impact; three broad themes, people, planet, and progress; and seven working groups of human capital, inclusion, trusted AI, resilience, science, democratizing AI resources, and social good. The summit highlighted use cases in agriculture, healthcare, education, and governance, and devoted considerable attention to the diffusion and adoption of AI technologies to the last mile.

What the Summit Delivered

The summit’s flagship diplomatic output was the India AI Impact Summit Declaration, endorsed by over ninety-two countries and international organizations, including China, Russia, the United Kingdom, the EU, and the United States.³ This constituted the broadest sign-on in the series by a significant margin. The declaration centers on voluntary, inclusive, and ethical AI governance, structured around the seven working groups. The signatories, spanning advanced economies and developing nations alike, endorsed a common framework oriented around development outcomes and equitable access. For India specifically, it validates the proposition that a Global South nation can set the terms of a major multilateral technology conversation, and that those terms can command broad international assent.

Other tangible outputs are the New Delhi Frontier AI Impact Commitments, signed by thirteen leading global and Indian frontier model developers, including Amazon, Meta, OpenAI, Anthropic, and Sarvam AI.⁴ These commitments centered on two core undertakings. The first focuses on transparency around real-world AI usage: participating firms agreed to analyze and publish anonymized, aggregated insights into how their AI systems are being used. This aims to help policymakers and researchers understand the impact of AI on jobs, skills, productivity, and economic transformation. The second focuses on inclusion, with companies committing to strengthening testing and evaluation of AI systems across underrepresented languages and cultural contexts to make frontier AI models more reliable and accessible.

The summit catalyzed over \$200 billion (₹18.6 lakh crore) in AI investment announcements, and India announced the expansion of its sovereign compute capacity with a target of 100,000 GPUs by the end of 2026.⁵

Opportunities Under “People, Planet, and Progress”

Under “People,” the summit elevated questions of AI literacy, workforce transition, and cultural and linguistic inclusion to a multilateral priority. The “Human Capital” and “Inclusion” working groups produced frameworks for equitable skilling and for AI systems that account for underrepresented communities.⁶ Embedding these concerns within the formal architecture of a global summit process gives them institutional weight that was lacking in earlier iterations.

Under “Planet,” the summit brought the environmental footprint of AI development into the governance mainstream. The “Resilience” working group addressed the energy demands of training and deploying AI models, a question of particular salience for countries with constrained infrastructure, and the guiding principles on resilient and efficient AI represent an early attempt to establish shared expectations around sustainability.⁷

Under “Progress,” the focus on democratizing AI resources, from compute and datasets to open-source models, directly addressed the structural asymmetries that determine which countries can participate in AI development as producers rather than consumers. The Charter for the Democratic Diffusion of AI, endorsed by twenty-five countries and international organizations, provides a voluntary framework for promoting access to foundational AI resources, supporting locally relevant innovation, and strengthening resilient AI ecosystems.⁸ The Global AI Impact Commons operationalizes the idea that cross-country learning on AI deployment should not remain ad hoc but should be institutionalized through a shared platform.⁹

The Way Forward

The India AI Impact Summit has set a precedent that will shape the trajectory of this multilateral process. By hosting and shaping the agenda, India demonstrated that the global AI conversation need not be defined exclusively by frontier model developers. The Global South framing introduces a different set of priorities into multilateral AI governance: not only how to govern AI safely, but how to ensure that governance frameworks do not entrench the existing concentration of technological capability. This provides a new direction that will be difficult to walk back as the process moves to Geneva in 2027.

The challenge ahead is to ensure that the frameworks and platforms launched in New Delhi are operationalized to deliver their stated objectives. The summit process has moved from safety to action to impact. The next phase will be determined by whether this impact can be sustained beyond the summit itself.

The AI Safety Journey From Bletchley to Delhi

Tejas Bharadwaj

Over the last three years, artificial intelligence (AI) has emerged as a critical subject of diplomacy through a series of high-level AI summits hosted by the United Kingdom, South Korea, France, and India. Discussions have focused on developing appropriate templates for AI governance, the sustainable development of AI infrastructure, identifying its impact and use cases for socio-economic development, and fostering inclusivity and diffusion. A common theme that ran through the summits was how stakeholders could facilitate AI innovation without compromising safety. The concept of AI safety as a core agenda has evolved through successive iterations of the summit.¹⁰ As it moves to Switzerland next year, this commentary will explore first how AI safety has featured across different summit processes and second how it can be taken forward at the next summit.

AI Safety Across the Summits

What began with a real sense of existential threats posed by frontier AI has gradually matured into discussions on making AI more trustworthy, secure, and reliable to maximize user benefits and societal impact.¹¹

The inaugural AI Safety Summit, convened at Bletchley in November 2023, saw the participation of twenty-eight countries, including the United States, China, and the EU, which agreed to the Bletchley Declaration.¹² The declaration established a consensus on the risks and opportunities associated with frontier AI systems and emphasized the need for increased collaboration on AI safety. Notably, countries and companies agreed to state-led testing of next-generation AI models before deployment, including through partnerships with state-backed AI Safety Institutes. The “Bletchley effect,” as it came to be described, led several countries, like the UK, Singapore, the EU, and the United States, to establish their own AI Safety Institutes.

Subsequently, at the AI Seoul Summit, co-hosted by the UK and South Korea in May 2024, the Seoul Declaration was ratified by Australia, Canada, France, Germany, Italy, Japan, Singapore, the UK, the United States, and the EU.¹³ The declaration, among other things, called for efforts to establish or expand AI safety institutes and promote cooperation on AI safety research through an international network. Notably, sixteen leading AI companies also agreed to the voluntary Frontier AI Safety Commitments.¹⁴

Thereafter, at the AI Action Summit in Paris in February 2025, the statement at the Paris AI Summit was endorsed by fifty-seven states, focusing on the impact of AI on sustainability, human rights, and inclusivity.¹⁵ Notably, the statement witnessed a shift in the AI safety rhetoric focused on existential threats that was emphasized in the past two summits.

The statement, however, did not see the United States and the UK join.¹⁶ The UK held back, noting that “the statement did not provide clarity on global governance nor go far enough to address the concerns around AI’s impact on national security.”¹⁷ While the United States’ position, derived from their vice president’s speech at the summit, indicated a need for discussions to shift toward opportunities that AI presents, noting that Europe’s excessive AI regulation could impede innovation.¹⁸ Days after the AI Action Summit, the UK AI Safety Institute, the first of its kind, was renamed the UK AI Security Institute, broadening its focus to include AI risks that have security implications.¹⁹ In June 2025, the United States, too, reformed its AI Safety Institute into the U.S. Center for AI Standards and Innovation.²⁰ Housed within the National Institute of Standards and Technology (NIST), it focuses on a more pro-innovative approach to AI governance, while according primacy to national security.²¹

In November 2025, the EU Commission proposed amendments to the AI Act under the “Digital Package” to provide the EU AI Office with a centralized oversight of AI systems in order to reduce fragmentation and bureaucratic bottlenecks while promoting innovation.²² These developments, undertaken since the AI Action Summit in Paris, reflect a shift in the approach of countries toward AI safety and the role of the early AI Safety Institutes. Previously, these institutes focused on AI risk management and the global coordination of safety metrics through international networks and bilateral data-sharing arrangements. The current AI Safety Institutes seek voluntary commitments from companies, actively participate in current global standard-setting forums, and test AI systems for national security risks (such as cybersecurity and biological threats), broadening their scope from earlier concerns regarding AI bias and freedom of speech.²³

A Shift From AI Safety to a More Secure and Trustworthy AI

In the lead-up to the AI Impact Summit, India moved beyond the discourse on frontier AI, its safety, and risks to unlock practical use cases (opportunities) and also gather perspectives from the Global South, with a focus on the democratization of AI resources.

The three pillars of the summit—people, planet, and progress—steered seven working groups, including one on “Safe and Trusted AI,” aimed at fostering transparency, reliability, and public trust in AI systems, while promoting responsible innovation and the equitable sharing of benefits.²⁴ The Delhi Declaration, endorsed by ninety-two countries, the EU, and the International Fund for Agricultural Development (IFAD), notably made no mention of the word “safety.”²⁵ Instead, a shift was seen toward secure and trusted AI, linking public trust to the unlocking of societal and economic benefits of AI across healthcare, agriculture, education, and public services. (See Table 1).

The declaration called for deepening the understanding of potential security aspects of AI, the importance of security in AI systems, industry-led voluntary measures, and the adoption of technical solutions, and appropriate policy frameworks that enable innovation while promoting the public interest throughout the AI’s life cycle.²⁶ The notable aspect of the summit was that the Delhi Declaration was endorsed by all G7 countries, along with China and Russia.²⁷ In particular, the endorsement by the UK and the United States marked a shift from the Paris declaration, which both had declined to sign. The endorsement signals India’s success in balancing divergent approaches to AI worldwide and indicates that unlocking AI-led innovation and its benefits is centered on fostering “public trust.” Furthermore, India broadened the earlier focus on safety in previous summits to address the security implications of AI.

Another outcome was the New Delhi Frontier AI Impact Commitments signed by global and Indian AI companies and included two voluntary commitments.²⁸ The first focused on promoting an understanding of the practical uses of AI through insights and evidence generated by these companies. The second called for strengthening the contextual effectiveness and evaluations of AI systems across cultures and languages, particularly underrepresented ones. Both commitments reflect the summit’s objective of promoting real world impact of AI in a contextual manner, across geographies, cultures, and languages. While “safety” did not feature in the summit declaration, the Safe and Trusted AI working group proposed the creation of an AI Safety Commons, a shared global repository for the testing, evaluation, and deployment of AI systems in the Global South.²⁹ This reflected India’s efforts to remain committed to the AI safety discussions, especially among countries in the Global South, despite not pushing for its inclusion in the summit declaration.

Table 1. How AI Safety has Featured Through the Summits

Summits	Reference to “AI Safety” in the declarations	Endorsement by Countries	Notable Inclusions or Omissions
AI Safety Summit, 2023 (Bletchley, UK)	15 references	28 countries and the EU	China, India, the EU, Japan, the UK, and the United States joined the commitment. Russia was a notable omission.
AI Seoul Summit, 2024 (Seoul, South Korea)	14 references	9 countries and the EU	Australia, Canada, the EU, France, Germany, Italy, Japan, South Korea, Singapore, the UK and the United States endorsed it. India, China, Russia, South Africa, and Brazil were notable omissions.
AI Action Summit, 2025 (Paris, France)	6 references	58 countries, the EU, and the African Union Commission	The United States and the UK were notable omissions.
AI Impact Summit, 2026 (Delhi, India)	No Mention of “Safety” Pivot to “Secure and Trusted AI.”	92 countries, the EU, and the IFAD	Russia, the UK, China, and the United States signed the declaration.

Sources: “The Bletchley Declaration by Countries Attending the AI Safety Summit, 1-2 November 2023,” Government of the United Kingdom, updated February 13, 2025, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>; “Seoul Declaration for safe, innovative and inclusive AI by participants attending the Leaders’ Session: AI Seoul Summit,” Government of the United Kingdom, May 21, 2024, <https://www.gov.uk/government/publications/seoul-declaration-for-safe-innovative-and-inclusive-ai-ai-seoul-summit-2024/seoul-declaration-for-safe-innovative-and-inclusive-ai-by-participants-attending-the-leaders-session-ai-seoul-summit-21-may-2024#contents>; “Statement on Inclusive and Sustainable Artificial Intelligence for People and the Planet,” Permanent Mission of France to the United Nations in New York, February 11, 2025, <https://onu.delegfrance.org/statement-on-inclusive-and-sustainable-artificial-intelligence-for-people-and>; “AI Impact Summit Declaration, New Delhi (February 18–19, 2026)” Ministry of External Affairs, Government of India, February 21, 2026, <https://www.mea.gov.in/bilateral-documents.htm?dtl/40809>.

What to Look Out for at the Next AI Summit

While the UN has aimed at fostering a dialogue for global AI governance, the current geopolitical appetite especially among countries at the forefront of AI innovation, to create a global watchdog on AI safety or security—similar to the International Atomic Energy Agency (IAEA) or the Organisation for the Prohibition of Chemical Weapons (OPCW)—is low.³⁰

An indication of this trend was the U.S. statement at the AI Impact summit, which underscored its strong disagreement with any mechanism aimed at global governance of AI, advocating instead for a mix of regulatory and voluntary frameworks under a broader national AI policy.³¹

Known for its neutrality and for hosting several major international and intergovernmental organizations, Switzerland finds itself at an opportune moment. Firstly, it can amplify what AI security or trustworthiness means for everyone, create norms that the aforementioned principles entail for compliance, and see if global capabilities can be built for these normative compliance. Second, it can initiate discussions on whether a CERN-like institution can be established to offer independent, scientific as well as contextual audits of AI models for the world, considering that evaluating AI models can be expensive and resource-heavy, and is often restricted to Big Tech firms or state-backed AI safety and security institutions.

The former requires Switzerland to leverage its experience as a hub of multilateral institutions to steer global normative conversations on technological security and test the appetite for global consensus on norms for diffusing security and trustworthiness in AI. The latter entails endorsing a global approach or institution that enables developing nations as well as startups to access evaluation tools and participate in the AI safety/security processes.

Conclusion

The journey through Bletchley, Seoul, Paris, and Delhi has helped build a foundational global consciousness on AI safety and shift toward AI security and trustworthiness. As we look toward the next AI summit, the challenge is no longer just to identify the risks but also to negotiate and test an international mechanism capable of addressing them.

Understanding the Global AI Impact Commons

Shruti Mittal and Nidhi Singh

AI is a transformative technology with the potential to redefine social and economic structures across the world. Individuals, companies, and countries are increasingly leveraging AI to improve public services, drive productivity, and address longstanding development challenges. Yet the narrative around AI remains primarily driven by a handful of technology companies concentrated in a few countries.³²

The 2026 AI Summit in India was guided by the mantra of “Impact,” shifting the conversation to a different set of questions: How is AI actually being used? To what effect? And where does the greatest potential for transformation lie? For countries across the Global South, these questions are the starting point for any serious engagement with AI as a tool for development.³³

The AI for Economic Growth and Social Good thematic working group was constituted under India’s AI Action Summit process to address this gap. Its mandate recognized that AI can advance development goals across healthcare, education, and agriculture, but that realizing this potential requires overcoming barriers related to capacity gaps, fragmented knowledge-sharing, and misaligned incentives. One of its principal outcomes, the Global AI Impact Commons,³⁴ translates this mandate into a practical instrument.³⁵ The Commons is a repository of high-quality impact stories from around the world that aims to provide access to AI resources and artefacts to enable cross-border learning and accelerate the development of AI solutions from pilots to population scale.

Why Build a Repository?

The global AI conversation has largely been dominated by discourse around large language models—their capabilities, compute requirements, and the competitive dynamics between the firms and nations building them. This is an important conversation, but not the most

relevant for the Global South. Countries in the Global South must find a middle ground that balances efforts to develop frontier models with increased efforts to deploy AI effectively within existing institutional and fiscal constraints. The Global AI Impact Commons recenters the conversation around use cases, focusing on what AI is doing on the ground for the people and systems it is meant to serve. The rationale is straightforward. Global South countries face structurally similar challenges, like overstretched public health systems, uneven access to education, vulnerable agricultural supply chains, and climate shocks. Where one country has developed an AI-enabled solution that works in such a context, others stand to benefit.

The summit, however, revealed that this cross-learning rarely happens systematically. Solutions remain siloed and knowledge dispersed. Even where promising AI solutions do exist, many remain trapped in pilot purgatory, successfully trialed but never scaled beyond initial deployments.³⁶ The repository is designed to break this pattern. By documenting the existing solutions and how they were deployed, sustained, and scaled, the repository provides the granular, operational detail needed to move from experimentation to implementation. It offers a structured mechanism for South-South knowledge transfer, grounded in what AI has done rather than what it could do.

What Makes the Commons Different?

The Commons is distinctive for three reasons. First, it focuses on global AI solutions that have already delivered measurable benefits to specific groups—farmers, students, and patients—rather than focusing on theoretical potential. It documents how these tools have improved lives in tangible ways. For example, by improving student learning, increasing diagnostic accuracy for clinicians, or reducing administrative backlogs for judicial officers. For policymakers in the Global South who are currently deciding how to allocate public resources, this proof of impact provides important evidence for procuring AI applications for sectoral use-cases in their local contexts.

Second, the Commons lays down a shared vocabulary for measuring how AI can deliver value. Without a shared vocabulary, AI success stories can remain anecdotal and appear as isolated examples rather than proven and scalable models. Standardizing how impact is reported could allow countries to make more meaningful comparisons between how AI solutions have fared across different geographical and linguistic contexts while evaluating their local relevance. Each use-case is documented in a standardized format that clearly lists its key outcomes, partnerships, and impact metrics. This consistent structure allows the user to compare different AI solutions that use the same technology but show varying results within and across sectors.

Third, the Commons collects data on the replicability and adaptability of the included AI solutions by detailing the technical, financial, and human resources required for a successful rollout. By documenting these specific requirements alongside key institutional partnerships,

the Commons allows other countries to move beyond a basic summary of a tool's function to understand its operational reality. For example, Adalat AI's deployment in Indian courtrooms is documented in the Commons, including its replicability and adaptation.³⁷ It highlights that Adalat AI's success is built on a shared common law history and modest resource requirements—factors that make the solution highly replicable in African contexts, where it is now being piloted. Because the story is framed through this shared vocabulary, policymakers can move beyond a standalone anecdote to see a proven and scalable model that is directly relevant to their own needs.

The Future of the Commons

The future of the Commons, however, depends on its ability to evolve from a static collection into a living resource. As is the case for all digital repositories, the primary friction will be maintenance. The Commons will need active, continuous curation to verify future submissions, update entries, and ensure consistency with evolving taxonomies in the AI landscape. It will also require a permanent institutional framework and stable financial backing to prevent it from becoming an unreliable archive over time. Dedicated support will enable the repository to scale alongside the growing number of AI solutions being deployed globally, ensuring it remains a reliable utility in the long term.

The Commons could also benefit from the intentional inclusion of “failure stories”—AI deployments that did not deliver their stated objectives, produced unintended user harm, or were simply discontinued. By documenting both successes and setbacks, the Commons could serve the dual purpose of democratizing best practices while demonstrating the self-examination required for responsible AI governance.

The Global AI Impact Commons is a product of its moment. It was born from the recognition that the countries that stand to gain the most from AI are the least represented in conversations about how it is built, governed, and shared. Yet the value of any commons lies not in its creation but in its use. Its promise will only be realized if the diversity of use cases from Global South countries is maintained and their metrics are actively updated. Whether it shapes outcomes or fades into irrelevance will depend not on the ambition of its mandate, but on the rigor of its implementation.

The Resource Cost of India's Data Center Expansion

Adarsh Ranjan

Introduction

Among the most significant outcomes of the India AI Impact Summit 2026 were investment announcements by leading technology companies to scale up the AI data center capacity in India.³⁸ Reliance Industries pledged \$110 billion over the next seven years; Adani Enterprises committed \$100 billion by 2035 to build renewable-powered AI data centers; and Tata Consultancy Services announced a partnership with OpenAI to build a 1 GW AI-ready data center, and Google announced a \$15 billion AI hub in Visakhapatnam.³⁹ These investments come on the back of India's push toward developing sovereignty across the AI stack, particularly infrastructure sovereignty, perceived to be the enabling layer for AI development and the backbone of digital infrastructure.⁴⁰

However, data centers are significant consumers of both electricity and water. The Economic Survey 2026 cautioned that “rapid expansion of artificial intelligence data centers can put severe stress on India's grid systems and groundwater capacity.”⁴¹ Amid the considerable enthusiasm accompanying the announcements, insufficient attention has been paid to the resource costs these facilities will impose and the safeguards required to manage them. As the scale of investment grows, the absence of binding sustainability requirements represents a significant gap in India's digital infrastructure strategy.

This article examines the resource pressures that data center expansion is likely to impose on India's water systems and electricity grids and argues that the current policy framework is ill-equipped to manage them.

Moreover, one of the outcomes of the AI Summit working groups was the Voluntary Guiding Principles on Resilient, Innovative, and Efficient AI, to which India is a signatory. It promoted “resource-conscious” and “energy-efficient” systems and supported the development of metrics designed to measure efficiency of AI, with a specific focus on

energy and water efficiency.⁴² This established not just a domestic policy imperative but an international commitment to a normative standard. As investments accelerate, this is a timely opportunity to integrate sustainability into the foundations of India's digital infrastructure.

Constrained Resources

Commitments to building data centers prior to and at the AI Summit reflect a broader strategic logic. India generates nearly 20 percent of global data but holds only around 3 percent of the total data center capacity.⁴³ While India's capacity had been growing steadily, the Union Budget 2026–2027 marked a strategic inflection point by introducing a long-term tax holiday for cloud providers using India-based infrastructure.⁴⁴ By reducing tax liabilities on global revenues and offering policy certainty, the measure reflected a clear shift from organic growth to state-led acceleration, signaling intent and setting the stage for some of the announcements that followed.

Meanwhile, India's emergence as a data center destination is anchored in its structural cost advantages. At \$7 per watt, India's data center development costs are among the lowest globally, alongside relatively low electricity tariffs.⁴⁵ These cost fundamentals significantly improve project viability for hyperscalers and underpin India's growing attractiveness as a competitive market for infrastructure. From a national perspective, such investments align with India's strategic objective of building domestic digital infrastructure to support AI capabilities, economic growth, and the expansion of the digital economy. They also advance data sovereignty by ensuring that critical data and compute capacity remain within the national jurisdiction, reducing dependence on external infrastructure, as mandated in the Digital Personal Data Protection (DPDP) Act, 2023.

Data centers are, however, resource-intensive by nature, and governments across the world are increasingly grappling with how to balance development with the infrastructure and environmental costs.⁴⁶

India's resource constraints make this balance particularly challenging. The NITI Aayog, the Indian government's public policy think tank, in its "Composite Water Management Index" report from 2019, noted that India was suffering from the most severe water crisis in its history, with almost 600 million people experiencing high to extreme water stress.⁴⁷ The situation is projected to worsen sharply, with water demand expected to be twice the available supply by 2030.

A data center consumes an estimated 25.5 million liters per megawatt (MW) every year. As India looks to increase its capacity to gigawatt scale, the estimated data center water consumption is expected to grow from an estimated 150 billion liters in 2025 to 358 billion liters by 2030.⁴⁸ This will only lead to an increase in the pressure on areas already facing acute water stress. If data centers depend heavily on groundwater or municipal sources, they may end up competing with agriculture and households.

The problem is geographical as much as aggregate. A majority of India's data centers are concentrated in Mumbai, Hyderabad, Delhi-NCR, Bengaluru, and Chennai, the most water-stressed cities in the country.⁴⁹ Additionally, a CEEW study shows that 57 percent of districts are at high to very high risk from extreme heat.⁵⁰ As a result, data centers located in these districts could also raise cooling loads substantially during summer months or aggravate peak-load management issues that may reinforce reliance on fossil-based generation.

Electricity demand is also expected to scale commensurately. Installed data center capacity in the country has quadrupled since 2020, growing from 375 MW to approximately 1.5 gigawatts (GW) in 2025.⁵¹ This is projected to grow further, from 1.5 GW to 9 GW by 2030.⁵² Consequently, electricity consumption from data centers is expected to rise from under 1 percent to approximately 3 percent of the total national usage over the same period.⁵³

The transmission and distribution infrastructure required to absorb this scale of demand does not currently exist and grid infrastructures will need considerable upgrades.⁵⁴ Hyperscale facilities can require close to a gigawatt of power, necessitating dedicated high-voltage grid substations to ensure uninterrupted power supply.⁵⁵ Without a national plan for adequate grid connection points, the result is suboptimal infrastructure additions, the costs of which could be passed on to consumers through higher tariffs.⁵⁶

This risk is exemplified by the state of India's distribution network, which recorded combined transmission and distribution losses of 16.64 percent of total electricity generated in FY2023–2024, which is more than double the global average of 6–8 percent.⁵⁷ Where reliable grid connections are unavailable, operators turn to captive generation. Data centers already routinely rely on diesel generators during grid failures. This further drives up emissions and costs to consumers as it fragments the demand on the grid from the data center, leaving transmission costs socialized across remaining consumers.

The Policy Gap

India does not have a national policy that governs the construction and operation of data centers in the country. Currently, a data center can be developed in India without having to undergo a comprehensive environmental impact assessment, as they are not explicitly covered under the schedules of the Environmental Impact Assessment Notification, 2006.⁵⁸ While the Ministry of Electronics and Information Technology's 2023 Green Data Center guidelines set benchmarks on power usage effectiveness (a metric used to assess how efficiently a data center uses energy), these remain advisory in nature and “carry no penal consequence” under the Environment Protection Act, limiting their enforceability in practice.⁵⁹

An earlier attempt to shape this space came in the form of the draft Data Centre Policy introduced by the Ministry of Electronics and Information Technology (MeitY) in 2020.⁶⁰ However, the policy was oriented primarily toward positioning India as a global data center hub, outlining measures to attract investment and streamline regulatory processes. The policy “encouraged” the use of clean and cost-effective power but set up no binding constraints on developers and operators and contained no interventions regulating the use of water. The draft was ultimately not formally notified.

In the absence of a central framework, states moved to fill the gap with their own data center policies. Karnataka, Telangana, Maharashtra, Gujarat, Uttar Pradesh, Odisha, Tamil Nadu, Haryana, and West Bengal are among the states that have enacted data center policies.⁶¹ Common incentives include land at concessional rates, capital subsidies, electricity tariff relief, stamp duty exemptions, and relaxations on building approvals. Akin to the 2020 Draft Data Center Policy, these policies are aimed at drawing investments into these states. Few mention sustainable practices at all, and those that do frame them as voluntary incentives rather than conditions of operation.

In August 2025, the MeitY circulated a draft of a revised policy on data centers, National Data Centre Policy 2025, and held consultations with industry representatives to gather their views.⁶² The draft itself has not been made publicly available, but statements from industry representatives offer some insight into the broad contours of the policy.⁶³ Much like its 2020 predecessor, the thrust of the proposed policy appears to be primarily oriented toward positioning India as a globally competitive destination for data center capital through fiscal incentives, simplified approvals, and enabling infrastructure. Questions remain about how the sustainable use of resources will be built into the policy.

Conclusion

India’s rapid expansion of its data center capacity is driven by clear strategic imperatives—advancing AI capabilities, securing data sovereignty, and closing the gap between its contribution to global data and its share of the global compute infrastructure. Yet this expansion is progressing against a backdrop of acute and worsening resource constraints. India’s water systems are already under extreme stress, and the cities where data centers are most concentrated are among the country’s most vulnerable. The country’s electricity infrastructure faces comparable pressures, with transmission and distribution losses far above global averages and grid capacity lagging behind projected demand. The result is a policy environment in which the resource costs of data center expansion remain largely unaddressed, even as India has committed, through international frameworks, to develop resource-conscious AI infrastructure.

India's AI Compute Question: Strategic Leverage Over Sovereignty Moonshots?

Konark Bhandari

For countries seeking a meaningful role in the current AI race, the stakes have never been higher. Scholars at the Carnegie Endowment for International Peace have argued that certain countries, including India, risk capturing the harms of AI while reaping minimal benefits, facing adversaries who might wield frontier AI against them.⁶⁴ With 1.4 billion people, a growing tech ecosystem, and \$10 billion in committed semiconductor investment, India sits at an inflection point.⁶⁵ The question is not whether AI matters to India's future—it unmistakably does—but whether the country's current strategy is enough to translate its size into genuine leverage over the compute infrastructure. This piece will look at how India can go about doing that.

Initial Steps

The Office of the Principal Scientific Adviser's (PSA) December 2025 white paper on democratizing AI infrastructure is a thoughtful document in this regard.⁶⁶ It proposes treating compute, datasets, and AI models as digital public goods, drawing on India's digital public infrastructure tradition. It recommends a phased approach starting with metadata standards and access protocols and gradually moving toward federated data access and coordinated compute exchange. But to answer the question concerning the geopolitics of compute, the PSA's framing is primarily domestic. It is a strategy for democratizing access within India, not for securing India's position in the global hierarchy of compute access. These are related but distinct problems.

As stated above, the proposed framework by certain scholars from Carnegie suggests that the more promising path is “leveraging bottlenecks.” This would entail finding a sector between AI progress and real-world effects and building economic and foreign policy around it, securing not just access to AI but a durable negotiating position with the great powers who control it. What, concretely, might India's bottlenecks be, and how can India build a durable AI infrastructure ecosystem?

Lessons from Other Sectors

The most obvious starting point is establishing the capability and capacity to manufacture hardware in general. India would need to comprehensively map all the layers of the AI hardware domain and work backward to identify where the most demand lies and how it can fulfill it. In this regard, the Indian government's gazette notification from the Ministry of Petroleum and Natural Gas is instructive.⁶⁷ It was notified amid heightened energy security concerns caused by the U.S.-Iran conflict that disrupted gas and liquefied petroleum gas supplies. On its surface, this was an emergency supply-chain measure that mandates all entities across the oil and gas value chain (refiners, liquefied natural gas importers, pipeline operators, and city gas distributors) to furnish granular data on production, imports, stocks, and consumption patterns.

The same data infrastructure, if extended thoughtfully, becomes a strategic asset of a different kind in the hardware sector. A government with real-time visibility into such production, storage, and consumption patterns is in a far better position to identify where data center capacity can actually be built, what headroom exists, and which industrial corridors are genuinely viable for AI infrastructure investment. If India can understand, in granular fashion, the demand outlook for items like chips, circuit boards, server racks, cooling systems, switchgears, and transformers, it could move from reactive to predictive in managing its AI infrastructure buildout.

Chips and Hardware—Integrating Into Pre-Existing Supply Chains

Then there is the chip supply chain question, where India has made tangible but underappreciated progress. Market experts have documented a marked inflection point in the role of central processing units (CPUs) in data centers: AI model training and inference are now using CPUs more intensively, and major firms like Intel saw an appreciable uptick in datacenter CPU demand in late 2025.⁶⁸ This may matter for India because CPU manufacturing, including outsourced assembly, testing, and packaging, is less technically demanding than frontier GPU fabrication. If CPUs continue their comeback as AI inference scales from training-heavy to inference-heavy workloads, India's emerging outsourced semiconductor assembly and test (OSAT) capacity could fit into exactly the right part of the supply chain at exactly the right moment.

The Apple connection also deserves particular attention here. Recent reporting has documented a fascinating shift in the Taiwan Semiconductor Manufacturing Company's (TSMC) customer dynamics.⁶⁹ TSMC's revenue from high-performance computing, which includes AI chips, surged 48 percent in 2025, while smartphone revenue grew at a comparatively modest 11 percent. Yet crucially, analysis suggests Apple's product lines span more than a dozen TSMC fabs, a wide breadth of manufacturing footprint that is distinct from Nvidia's needs. India has already embedded itself in Apple's manufacturing

story—it accounts for \$50 billion in exports from India.⁷⁰ The question is whether India can move beyond being an assembly hub for Apple’s consumer devices and position itself as a supply-chain partner for the broader chip portfolio that keeps Apple relevant to TSMC. The India Semiconductor Mission (ISM) 2.0, announced in India’s budget in 2026 with a focus on semiconductor equipment and materials, design IP, and supply chain fortification, is the right instinct.⁷¹ But merging its ambitions more explicitly with the Production Linked Incentive Scheme for electronics succession that the government is actively redesigning with a greater emphasis on value addition and export performance would create a more integrated industrial strategy than currently exists.

This brings India to the trade competitiveness question that no AI strategy can avoid. India’s share in the global economy has remained almost stagnant over the past decade, reflected its limited participation in global value chains. Being embedded in AI hardware supply chains would need to go hand-in-hand with an export-strong economy. Indeed, this is likely how any leverage may be exercised by India. The AI hardware opportunity gives India a reason to revisit that reluctance, because the sectoral case for input liberalization in electronics is now undeniable, with India also signing a spate of free trade agreements (FTA) lately to encourage the same.

But all of this must be weighed against the legitimate concerns about the sustainability of the current AI investment cycle. The capital expenditure math is daunting: The United States’ AI data center capital expenditure for 2025 alone was \$400 billion.⁷² In India as well, experts have estimated that the AI data center boom will significantly boost capital expenditure.⁷³ Inventory costs, construction timelines, energy constraints, and water stresses are real risks that attach to any country betting heavily on physical AI infrastructure. Though an immediate return is not always guaranteed, the risk of a correction in AI investment must not deter India from engaging in the infrastructure story. It is a reason to be strategic about which parts of the infrastructure story to own.

Conclusion

The negotiating terms for India may worsen as frontier AI gets bigger. Acting early, building supply-chain leverage, and accumulating data assets now may cost far less than trying to negotiate from a weaker position in the future. Signing the Pax Silica was a good start. The next chapter needs to be about how India converts domestic capability into an international negotiating position.

Rethinking Inclusion in AI

Charukeshi Bhatt

As the first in the series of global AI summits held in the Global South, the India AI Impact Summit placed much needed emphasis on how the benefits of AI translate into the real world. In doing so, the summit championed diversity, inclusivity, and context-specificity in AI systems. The summit addressed the need for multilingual diversity more explicitly than any of the previous convenings, going beyond the broad policy language of earlier summits, which spoke of inclusivity largely in terms of accessibility, partnership, and multistakeholder process.⁷⁴ In those conversations, inclusivity had more to do with the number of stakeholders involved than the stakeholders themselves.

This distinction matters greatly for the Global South. In the discourse on AI, datasets occupy a foundational position in conversations about fairness. They are widely treated as the point of origin, the place where bias either enters or is corrected.⁷⁵ However, datasets are end products in themselves, the visible final output of a deeper pipeline that is rarely examined.

This piece argues that the inclusion debate must move upstream, from viewing datasets as merely the input for model training to thinking about datasets as a collective output emerging from an iterative and participatory process. The inclusion debate must engage with questions that predate the creation of datasets: Who produces language? Who digitizes it? And who participates in data creation? This upstream move rests on how the dynamic between data, language, and communities vis-à-vis AI models is conceptualized. A two-way rethinking of this relationship is necessary.

Why Languages Matter for AI

First, at a broader level, we need a new way of thinking about languages in AI systems. The view that including more languages during model training is mainly about representation needs to change. Instead, training AI models on a wide range of languages is essential to expand their usefulness and ensure they work effectively across different contexts. Their usage in sectors like healthcare, education, agriculture, and climate depends on how well

these models respond to queries generated in the users' language. Consider this: An image generation model, when asked multiple times to reproduce a reference image exactly as it is, without any changes, produces a completely different image each time. This is because the model reproduces patterns with which it has most dominantly been acquainted through its training data. So, when it encounters a visual image with unrepresented elements, it tries to fill in the gap with what it can borrow from that data. The result is a completely different image every time.

The parallel for language models is direct and worrying. Being trained in a limited set of languages, a model misses out on the meanings, grammar, logic, and reasoning of completely different knowledge systems. When such models receive queries in underrepresented languages, they produce answers that may be fluent but borrowed from a different logic system. While the drift in meaning is serious for factual queries, for culturally specific ones—such as those relating to agriculture or local law—the consequences may be far more serious.

While the drift can be visible for image models, it could go unnoticed in language models, especially for populations without the means for verification. The move upstream to issues of language systems is, therefore, not a simple question of fairness or representation, but of model efficacy and quality. Both these metrics impact the scale of adoption of an AI solution. Non-representative datasets will build models that are completely unexposed to certain structures of meaning, grammar, logic, and reasoning, resorting to compensatory behavior on the basis of training data when posed with very specific queries, leading to lower efficacy, trust, and adoption. In this sense, investing in inclusive data is not only a social imperative, but also a technical one. Recognizing its value for improving performance can also push model and solution developers to invest in the systems that make such data possible. This brings us to the second rethink.

Why Community Matters for Languages

The second shift is about identifying where the process of building inclusive AI begins. Languages are living resources—fluid, multicultural, and community-held. Making a language AI-ready is not simply a matter of digitizing a corpus. Many languages in the Global South present real technical challenges: different tonality, multiple or absent writing systems, and agglutinative structures. These factors challenge simple digital conversion.

Further, tonal nuances are particularly important as a number of AI use cases or solutions will be voice-based. As large language models (LLMs) increasingly shape how people communicate and make meaning, a model trained on an inadequate version of a language risks displacing the original.⁷⁶ Ethnographic research from Africa highlights how digitized language, drawn largely from formal sources like government documents and legislation, felt alien to native speakers in everyday use.⁷⁷ This is not a minor mismatch. In many Global South contexts, the gap between formal and vernacular languages carries a post-colonial weight, and an LLM that speaks only in a formal register may alienate the very communities it claims to serve.

For LLMs to be genuinely useful to farmers, small traders, or primary school teachers, they must be trained on language as it is spoken and used every day, built through sustained engagement with local speakers, linguists, sociologists, and ethnographers. Inclusion at this level involves constructing the linguistic and community infrastructure that creates datasets in the first place. The starting point is not the dataset itself, but addressing the structural absence of communities from the process that creates it. This would mean setting up local language centers where speakers, teachers, and domain experts work together to document everyday language use, including idioms, oral knowledge, and context-specific meanings. The Masakhane project in Africa is using a community-led approach to build datasets with participation from researchers, academics, and native speakers.⁷⁸ Such work will need to be continuous rather than being limited to one-time data collection and will require greater funding support for linguistic research. It will also mean building long-term, paid systems for community members to contribute to data annotation and validation, treating language work as skilled labor. However, safeguards are needed to ensure that such efforts do not expose Global South communities to new forms of harm, including visually disturbing content.⁷⁹ More broadly, building linguistic ecosystems will require moving beyond data scraping to work directly with farmers, traders, health workers, and educators to develop context-specific vocabularies.⁸⁰ This will address gaps stemming from missing documentation of local knowledge systems.

Investing Upstream in Language and Community Infrastructure

Some efforts from the Global South demonstrate what is possible. India's Bhashini initiative has built open language datasets across dozens of Indian languages through community participation under the Bhasha Daan initiative, while Africa's Masakhane project has brought together researchers and native speakers to develop natural language processing (NLP) tools for African languages largely absent from mainstream AI.⁸¹

At the AI Impact Summit, leading AI companies also committed to strengthening multilingual and use-case evaluations to better reflect diverse cultural contexts.⁸² These are important steps but taking this further will require a shift from focusing only on datasets to the systems and conditions that enable them.

At the same time, inclusivity cannot be treated as a responsibility of the Global South alone. Representative data is not just about fairness. It is essential for building effective AI systems that can support meaningful use cases. Model developers, then, will need to play a more active role in supporting and sustaining these linguistic ecosystems. Without this, efforts at inclusion will remain limited. Building inclusive AI is, in the end, as much about the people behind the data as it is about the data itself.

About the Authors

Konark Bhandari is a fellow with Carnegie India. He is a lawyer who has researched certain areas in the digital economy, focusing primarily on approaches to antitrust regulation of companies in the digital realm. He is also an expert on the regulation of the space technology ecosystem in India and has also written on India's approach to international space partnerships through bilateral and multilateral accords. Konark contributes regularly to discussions on the semiconductor ecosystem.

Tejas Bharadwaj is a senior research analyst with the Technology and Society Program at Carnegie India. He works on space law and policies, tracking India's space sector developments as well as issues pertaining to space security and sustainability globally. He also works on AI in military domain, including Lethal Autonomous Weapon Systems (LAWS), defense tech partnerships and cybersecurity policies.

Charukeshi Bhatt is a research analyst at Carnegie India, where her work focuses on the intersection of emerging technologies and international security. Her current research explores how advancements in technologies such as AI are shaping global disarmament frameworks and security norms.

Shruti Mittal is a research analyst in the Technology and Society Program at Carnegie India. Her work focuses on semiconductor industrial policy, AI governance, and open approaches to development for the Global South.

Adarsh Ranjan is a research analyst at Carnegie India where his research focuses on AI and emerging technologies, digital transformation, and technology partnerships. His current research explores India's evolving policy on AI compute and digital transformation in Global South countries.

Nidhi Singh is an associate fellow at Carnegie India. Her current research interests include data governance, artificial intelligence and emerging technologies. Her work focuses on the implications of information technology law and policy from a Global Majority and Asian perspective.

Notes

- 1 Ministry of Electronics & IT, “India AI Impact Summit 2026 Witnessed Extensive Participation of Public, with approximately 6 Lakh Attendees in Person,” press release, Press Information Bureau, March 11, 2026, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2238195®=3&lang=2>.
- 2 Ibid.
- 3 Ibid.
- 4 “New Delhi Frontier AI Impact Commitments, India AI Impact Summit 2026,” India AI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/New_Delhi_Frontier_AI_Impact_Commitments_f81b331b93.pdf.
- 5 Ministry of Electronics & IT, “India AI Impact Summit 2026: Landmark Global Declaration and Major AI Investment Commitments;” Soumyarendra Barik, “India’s GPU Capacity May Triple to 100,000 by 2026-end: IndiaAI Mission CEO,” *The Indian Express*, February 7, 2026, <https://indianexpress.com/article/business/indiaai-mission-ceo-gpu-capacity-triple-100000-2026-end-10519763/>.
- 6 “Human Capital Working Group Voluntary Guiding Principles for Skilling and Reskilling in the Age of AI,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/Guiding_Principles_for_Skilling_and_Reskilling_in_the_Age_of_AI_5a1c00d6a8.pdf; “Inclusion for Social Empowerment Working Group, Alliance for Advancing Inclusion Through AI,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/Alliance_for_Advancing_Inclusion_Through_AI_34e907558e.pdf.
- 7 “Resilience, Innovation and Efficiency Working Group, Voluntary Guiding Principles on Resilient, Innovative, and Efficient AI,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/Guiding_Principles_on_Resilient_Innovative_and_Efficient_AI_d527829f89.pdf.
- 8 “Democratizing AI Resources Working Group, Charter for the Democratic Diffusion of AI,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/Charter_for_the_Democratic_Diffusion_of_AI_a273df0545.pdf.
- 9 “About,” AI Impact Commons, accessed April 27, 2026, <https://www.aiimpactcommons.global/about>.
- 10 Cameron F. Kerry and Elham Tabassi, “Sovereignty, safety, and scale: Takeaways from the India AI Impact Summit,” Brookings Institution, March 12, 2026, <https://www.brookings.edu/articles/takeaways-from-the-india-ai-impact-summit/>.
- 11 James Gregory and Zoe Kleinman, “Rishi Sunak says AI has threats and risks - but outlines its potential,” *BBC*, October 26, 2023, <https://www.bbc.com/news/uk-67225158>.
- 12 “The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023,” Government of the United Kingdom, updated February 13, 2025, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>.
- 13 “Seoul Declaration for Safe, Innovative and Inclusive AI by Participants Attending the Leaders’ Session: AI Seoul Summit, 21 May 2024,” Government of the United Kingdom, May 21, 2024, <https://www.gov.uk/government/publications/seoul-declaration-for-safe-innovative-and-inclusive-ai-ai-seoul-summit-2024/seoul-declaration-for-safe-innovative-and-inclusive-ai-by-participants-attending-the-leaders-session-ai-seoul-summit-21-may-2024>.

- 14 “Frontier AI Safety Commitments: AI Seoul Summit 2024,” Government of the United Kingdom, updated February 7, 2025, <https://www.gov.uk/government/publications/frontier-ai-safety-commitments-ai-seoul-summit-2024/frontier-ai-safety-commitments-ai-seoul-summit-2024>.
- 15 “Statement on Inclusive and Sustainable Artificial Intelligence for People and the Planet,” Permanent Mission of France to the United Nations in New York, February 11, 2025, <https://onu.delegfrance.org/statement-on-inclusive-and-sustainable-artificial-intelligence-for-people-and-planet>.
- 16 Ibid.
- 17 Dan Milmo and Eleni Courea, “US and UK refuse to sign Paris summit declaration on ‘inclusive AI,” *The Guardian*, February 11, 2025, <https://www.theguardian.com/technology/2025/feb/11/us-uk-paris-ai-summit-artificial-intelligence-declaration>.
- 18 “VP JD Vance on the future of artificial intelligence,” YouTube video posted by “LiveNOW from FOX,” February 11, 2025, accessed April 28, 2026, <https://www.youtube.com/watch?v=64E9O1Gv99o>.
- 19 Department for Science, Innovation and Technology, AI Security Institute, and The Rt Hon Peter Kyle MP, “Tackling AI Security Risks to Unleash Growth and Deliver Plan for Change,” press release, Government of the United Kingdom, February 14, 2025, <https://www.gov.uk/government/news/tackling-ai-security-risks-to-unleash-growth-and-deliver-plan-for-change>.
- 20 “Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation,” press release, U.S. Department of Commerce, June 3, 2025, <https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai>.
- 21 Ibid.
- 22 “An agile Digital Rulebook for the EU,” European Commission, accessed April 28, 2026, <https://digital-strategy.ec.europa.eu/en/policies/digital-rulebook>.
- 23 “Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation,” press release, U.S. Department of Commerce.
- 24 Prime Minister’s Office, “PM to inaugurate India AI Impact Summit 2026 on 19 February,” press release, Press Information Bureau, February 18, 2026, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2229611>; Ministry of Electronics & IT, “IIT Madras and IndiaAI Mission to Host Global Conclave on Safe & Trusted AI Ahead of India - AI Impact Summit 2026,” press release, Press Information Bureau, December 9, 2025, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2200799>.
- 25 “AI Impact Summit Declaration, New Delhi (February 18–19, 2026),” Ministry of External Affairs, Government of India, February 21, 2026, <https://www.mea.gov.in/bilateral-documents.htm?dtl/40809>; IndiaAI, (@OfficialINDIAai), X post, March 9, 2026, <https://x.com/OfficialINDIAai/status/2030969349117296794>.
- 26 “AI Impact Summit Declaration, New Delhi (February 18–19, 2026),” Ministry of External Affairs, Government of India.
- 27 Ministry of Electronics & IT, “India AI Impact Summit 2026: Landmark Global Declaration and Major AI Investment Commitments,” press release, Press Information Bureau, March 2, 2026, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2234343>.
- 28 Ministry of Electronics & IT, “Championing Inclusive and Multilingual AI for the Global South, India Unveils New Delhi Frontier AI Commitments,” press release, Press Information Bureau, February 19, 2026, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2230201/>.
- 29 Ministry of Electronics & IT, “IIT Madras and IndiaAI Mission to Host Global Conclave on Safe & Trusted AI Ahead of India - AI Impact Summit 2026,” press release, Press Information Bureau, December 9, 2025, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2200799&lang=1®=3>.
- 30 “Global Dialogue on Artificial Intelligence Governance, 6-7 July 2026, Geneva,” United Nations, accessed April 28, 2026, https://www.un.org/global-dialogue-ai-governance/sites/default/files/2026-04/draft_note_on_themes_and_structure.pdf.

- 31 “U.S. Promotes AI Adoption, Sovereignty, and Exports at India AI Impact Summit,” The White House, February 20, 2026, <https://www.whitehouse.gov/releases/2026/02/u-s-promotes-ai-adoption-sovereignty-and-exports-at-india-ai-impact-summit/>.
- 32 Kate Brennan, Amba Kak, and Sarah Myers West, “Chapter 2: Heads I Win, Tails You Lose: How Tech Companies Have Rigged the AI Market,” in *Artificial Power: AI Now 2025 Landscape* (AI Now Institute, June 3, 2025), <https://ainowinstitute.org/2025-landscape>.
- 33 Narendra Modi, “PM’s Address at India AI Impact Summit 2026,” speech, Prime Minister of India, February 19, 2026, https://www.pmindia.gov.in/en/news_updates/pms-address-at-india-ai-impact-summit-2026/.
- 34 Website homepage, “Global AI Impact Commons,” accessed April 27, 2026, <https://www.aiimpactcommons.global/>.
- 35 Carnegie India was a knowledge partner to the AI for Economic Growth and Social Good working group and a contributor to the Global AI Impact Commons.
- 36 Shalini Kapoor and Tanvi Lall, “AI Adoption Journey for Population Scale,” Carnegie India, January 12, 2026, <https://carnegieindia.org/india/posts/2026/01/ai-adoption-journey-for-population-scale>.
- 37 “AI Powered Courtroom Transcription and Workflow Automation for Timely Justice” Global AI Impact Commons, accessed April 27, 2026, <https://www.aiimpactcommons.global/impact-stories/a1babd1a-3752-4ab8-84cf-c3c5be5ca1df?q=adalat+AI>.
- 38 Press Trust of India, “India AI Impact Summit 2026: \$240 Billion Investment from Global Leaders” *ETCIO.com*, March 3, 2026, <https://cio.economictimes.indiatimes.com/news/artificial-intelligence/india-ai-impact-summit-2026-240-billion-investment-from-global-leaders/128958486>.
- 39 Chandini Monnappa, Surbhi Misra and Urvi Manoj Dugar, “Reliance, Adani drive India’s AI push with plans to invest \$210 billion,” Reuters, February 19, 2026, <https://www.reuters.com/world/india/ambanis-reliance-will-invest-110-billion-ai-2026-02-19/>; “N Chandrasekaran Unveils Tata’s 1GW AI Data Centre, Chip & AI OS Plan,” YouTube video, by “Business Today,” February 19, 2026, https://youtu.be/_5ktbC1iaHc?t=326; Thomas Kurian, “Our First AI Hub in India, Powered by a \$15 Billion Investment,” *India (blog)*, Google, October 14, 2025, <https://blog.google/intl/en-in/company-news/our-first-ai-hub-in-india-powered-by-a-15-billion-investment/>.
- 40 “Eco Survey Flags Strain on Grids, Groundwater Due to AI Data Centres,” *Business Standard*, January 29, 2026, https://www.business-standard.com/budget/news/eco-survey-flags-strain-on-grids-groundwater-due-to-ai-data-centres-126012901461_1.html.
- 41 Animesh Jain, “Democratizing Access to AI Infrastructure,” *India’s AI Policy Priorities White paper Series*, (Office of the Principal Scientific Adviser to the Government of India, December 2025), https://psa.gov.in/CMS/web/sites/default/files/publication/WP_Democratising%20Access_V3.0_29122025A.pdf.
- 42 “Voluntary Guiding Principles on Resilient, Innovative, and Efficient AI,” India AI, February 2026, https://d19ob9sqt2wc.cloudfront.net/stage/uploads/Guiding_Principles_on_Resilient_Innovative_and_Efficient_AI_d527829f89.pdf.
- 43 *Attracting AI Data Centre Infrastructure Investment in India*, (Deloitte, May 2025), <https://www.deloitte.com/content/dam/assets-zone1/in/en/docs/services/engineering-ai-data/in-eaid-accelerating-ai-infrastructure-investments-in-india.pdf>.
- 44 “Budget 2026–27 Sets the Stage for India as a Global Hub for Cloud and AI Infrastructure,” press release, Press Bureau of India, February 14, 2026, <https://www.pib.gov.in/PressReleaseDetailm.aspx?PRID=2227953>.
- 45 Press Trust of India, “Data centre development cost in India at \$7 per Watt among global lowest, shows strong growth potential: Report,” *The Hindu Businessline*, September 16, 2025, <https://www.thehindubusinessline.com/info-tech/data-centre-development-cost-in-india-at-7-per-watt-among-global-lowest-shows-strong-growth-potential-report/article70055760.ece>.
- 46 Lucas Machado, Julia Ezcurdia, and Michaela Karamperi, “Policy Spotlight: Data Centres, AI Expansion, and the Environmental Impacts Governments Face,” *Fiscal Note (blog)*, August 25, 2025, <https://fiscalnote.com/blog/ai-data-centres-global-energy-water>.

- 47 Alina Vaidya Mahadevan, "A Parched Nation: Analyzing India's Water Scarcity Challenges," *Economic and Political Weekly* Vol. 59, Issue No. 20, May 18, 2024, <https://www.epw.in/engage/article/parched-nation-analyzing-indias-water-scarcity>; Press Trust of India, "By 2030, India's water demand to be twice the available supply indicating severe water scarcity: Report," *The Economic Times*, June 21, 2018, <https://economictimes.indiatimes.com/news/economy/agriculture/by-2030-indias-water-demand-to-be-twice-the-available-supply-indicating-severe-water-scarcity-report/articleshow/64679218.cms?>
- 48 Anwesha Sen, *Building India's Data Centres*, (Takshashila Institution, October 2025), <https://takshashila.org.in/content/publications/assets/20251029-Building-Indias-Data-Centres.pdf>; Nikhil Inamdar, "India's data centre boom confronts a looming water challenge," *BBC*, November 10, 2025, <https://www.bbc.com/news/articles/cgr417pwek7o>.
- 49 Kumar Gandharv, "India's data centre capacity crosses 1.5 GW as metro hubs drive growth," *CRNASIA*, March 19, 2026, <https://www.crnasia.com/india/news/2026/india-s-data-centre-capacity-crosses-1-5-gw-as-metro-hubs-drive-growth>.
- 50 Aman Malik and Sabarish Elango, "Why India must plan for AI data centres' energy and water footprint," Council on Energy, Environment and Water (blog), January 19, 2026, <https://www.ceew.in/blogs/why-is-water-based-cooling-a-big-issue-for-ai-data-centres-in-india>.
- 51 Ministry of Electronics & IT, "Data centre capacity in the country has increased from about 375 MW in 2020 to around 1500 MW by 2025," press release, Press Bureau of India, March 13, 2026, <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2239616>.
- 52 "Data centre boom: \$50 billion push to 9GW capacity by 2030," *ETEnterpriseAi.com*, October 22, 2025, <https://enterpriseai.economictimes.indiatimes.com/amp/news/industry/indias-data-centre-industry-set-for-50-billion-investment-and-9gw-capacity-by-2030/124735180>.
- 53 Soon Chen Kang and Ankita Chauhan, "Will India become a leading global datacenter market?," S&P Global, September 17, 2025, <https://www.spglobal.com/en/research-insights/special-reports/india-forward/shifting-horizons/will-data-center-growth-india-propel-country-global-hub-status>.
- 54 Press Trust of India, "India's data centre hub potential: Power, grid challenges and renewable integration, says Deloitte", *ETEnergyworld.com*, February 22, 2026, <https://energy.economictimes.indiatimes.com/news/power/india-emerges-as-a-data-centre-hub-challenges-and-opportunities-for-sustainable-ai-infrastructure/128670297>.
- 55 Pratyush Deep, "Can India's power system keep up with the explosion in data centres?," *The Indian Express*, February 18, 2026, <https://indianexpress.com/article/explained/explained-economics/india-power-system-data-centre-boom-10537422/>; Devansh Mittal, "Inside India's giant AI 'refineries': The high-security monoliths powering your ChatGPT prompts," *The Indian Express*, March 18, 2026, <https://indianexpress.com/article/long-reads/the-buildings-that-power-ai-inside-an-ai-data-centre-10587670/>.
- 56 Josh Saul et al., "AI Data Centers Are Sending Power Bills Soaring," *Bloomberg*, September 30, 2025, <https://www.bloomberg.com/graphics/2025-ai-data-centers-electricity-prices/>.
- 57 Anwesha Sen, *Building India's Data Centres*, (Takshashila Institution, October 2025).
- 58 Hera Rizwan, "Interview: In India, A Data Centre Can Be Built Without Any Environmental Assessment," *BOOM*, February 27, 2026, <https://www.boomlive.in/decode/interview-in-india-a-data-centre-can-be-built-without-any-environmental-assessment-30754?>
- 59 Ibid.
- 60 "Draft Data Centre Policy 2020," Ministry of Electronics & Information Technology (e-Governance Division), https://www.nitiforstates.gov.in/public-assets/Policy/policy_files/PNC510C000384.pdf.
- 61 Alenjith K. Johny and Puran Choudhary, "Execution tops sops as states race to attract data centres," *The Economic Times*, January 1, 2026, <https://economictimes.indiatimes.com/tech/startups/execution-tops-sops-as-states-race-to-attract-data-centres/articleshow/126277565.cms?>
- 62 "Industry Consultation Meeting of National Data Centre Policy 2025," National e-Governance Division, accessed April 28, 2026, <https://negd.gov.in/gallery/industry-consultation-meeting-of-national-data-centre-policy-2025/>.

- 63 Deborah Grey, “SPECIAL FEATURE | India’s Draft National Data Center Policy 2025: Industry leaders cautiously optimistic,” *WMedia*, September 21, 2025, <https://w.media/indias-draft-national-data-center-policy-2025-industry-leaders-cautiously-optimistic/>.
- 64 Anton Leicht, “What AI Summits Are For,” *Threading the Needle (blog)*, February 16, 2026, <https://writing.antonleicht.me/p/what-ai-summits-are-for>.
- 65 “Centre Commits Entire \$10 Billion Under ISM, Second Edition to Be Released Soon: S Krishnan,” *ETTelecom.Com*, September 3, 2025, <https://telecom.economicstimes.indiatimes.com/news/policy/india-semiconductor-mission-secures-10-billion-commitment-phase-two-announcement-imminent/123651135>.
- 66 Animesh Jain, “Democratizing Access to AI Infrastructure,” *India’s AI Policy Priorities White paper Series*, (Office of the Principal Scientific Adviser to the Government of India, December 2025), https://psa.gov.in/CMS/web/sites/default/files/publication/WP_Democratising%20Access_V3.0_29122025A.pdf.
- 67 Press Trust of India, “Govt classifies energy data as national security matter; mandates disclosure across oil, gas chain,” *The Pioneer*, March 20, 2026, <https://dailypioneer.com/news/govt-classifies-energy-data-as-national-security-matter-mandates-disclosure-across-oil-gas-chain>.
- 68 “Intel Reports Fourth-Quarter and Full-Year 2025 Financial Results,” Intel, press release, January 22, 2026, <https://www.intc.com/news-events/press-releases/detail/1759/intel-reports-fourth-quarter-and-full-year-2025-financial>; Gerald Wong and Dylan Patel, “CPUs are Back: The Datacenter CPU Landscape in 2026,” *SemiAnalysis*, Substack, February 9, 2026, <https://newsletter.semianalysis.com/p/cpus-are-back-the-datacenter-cpu>.
- 69 Tim Culpam, “Apple is Fighting for TSMC Capacity as Nvidia Takes Center Stage,” *Culpium*, Substack, January 15, 2025, <https://www.culpium.com/p/exclusiveapple-is-fighting-for-tsmc>.
- 70 Kiran Rathee, “Apple AmPLified! India Ships Out iPhones Worth \$50 Billion Till December 2025,” *The Economic Times*, January 5, 2026, <https://economicstimes.indiatimes.com/industry/cons-products/electronics/apple-amplified-india-ships-out-iphones-worth-50-billion-till-december-2025/articleshow/126339797.cms?from=mdr>.
- 71 Website homepage, India Semiconductor Mission, accessed April 27, 2026, <https://ism.gov.in/>.
- 72 Dwarkesh Patel and Romeo Dean, “Thoughts on the AI Buildout,” *Dwarkesh Podcast*, Substack, October 22, 2025, <https://www.dwarkesh.com/p/thoughts-on-the-ai-buildout>.
- 73 Alpha Desk, “AI Data Centre Boom May Lift Capital Goods Orders in India: IIFL Names Netweb, ABB Among Picks,” *CNBCTV18*, February 23, 2026, <https://www.cnbctv18.com/market/stocks/ai-capex-cycle-india-renu-baid-netweb-cummins-lt-abb-apar-hitachi-cgpower-mtar-iifl-data-centre-alpha-article-19855604.htm>.
- 74 “New Delhi Frontier AI Impact Commitments, India AI Impact Summit 2026,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloudfront.net/stage/uploads/New_Delhi_Frontier_AI_Impact_Commitments_f81b331b93.pdf; “The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023,” Government of the United Kingdom, November 1, 2023, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>.
- 75 Emily Diana, “Building AI Fairness by Reducing Algorithmic Bias,” *Tepperspectives*, Carnegie Mellon University, Tepper School of Business, July 18, 2025, <https://tepperspectives.cmu.edu/all-articles/building-ai-fairness-by-reducing-algorithmic-bias/>.
- 76 “How People Are Using ChatGPT,” OpenAI, September 15, 2025, <https://openai.com/index/how-people-are-using-chatgpt/>.
- 77 Lindsey DeWitt Prat, et al., “Decolonizing LLMs: An Ethnographic Framework for AI in African Contexts,” *EPIC Proceedings* (2024), 45–84, <https://www.epicpeople.org/decolonizing-llms-ethnographic-framework-for-ai-in-african-contexts/>.
- 78 “Projects,” Masakhane, accessed April 28, 2026, <https://www.masakhane.io/masakhane-african-languages-hub/projects>.

- 79 Chris Vallance, “Regulator contacts Meta over workers watching intimate AI glasses videos,” *BBC*, March 4, 2026, <https://www.bbc.com/news/articles/c0q33nvj0qpo>.
- 80 AIM, LinkedIn post, April 28, 2026, https://www.linkedin.com/posts/analytics-india-magazine_ekacare-nvidiaai-medicalai-activity-7429844635931328512-aOMs/.
- 81 “BhashaDaan,” Digital India Bhashini Division, accessed April 28, 2026, <https://bhashadaan.bhashini.co.in/bhashadaan/en/home>.
- 82 “New Delhi Frontier AI Impact Commitments,” IndiaAI, February 2026, https://d19ob9sqegt2wc.cloud-front.net/stage/uploads/New_Delhi_Frontier_AI_Impact_Commitments_f81b331b93.pdf.

Carnegie India

Carnegie India, established in 2016 and based in New Delhi, is part of a robust global network that includes more than 200 scholars. The center focuses primarily on three interrelated topics: technology and society, political economy, and security studies. Led by Indian experts with decades of policy experience, Carnegie India engages with governments, policymakers, academics, students, industries, practitioners, and civil society to provide insightful and fresh analysis on India's pressing challenges and its rising global role.

Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 200 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.



CARNEGIE INDIA

Carnegie India
Unit C-4, C-5 & C-6, Edenpark,
Shaheed Jeet Singh Marg
New Delhi – 110016, India

CarnegieIndia.org