CARNEGIE INDIA

Policy Planning & Research Division
Ministry of External Affairs
Government of India

2025

GLOBAL TECHNOLOGY SUMMIT

AI IMPACT SUMMIT
भारत 2026 INDIA

सर्वजन हिताय | सर्वजन सुखाय
WELFARE FOR ALL | HAPPINESS OF ALL

INNOVATION DIALOGUE

December 10-11, 2025

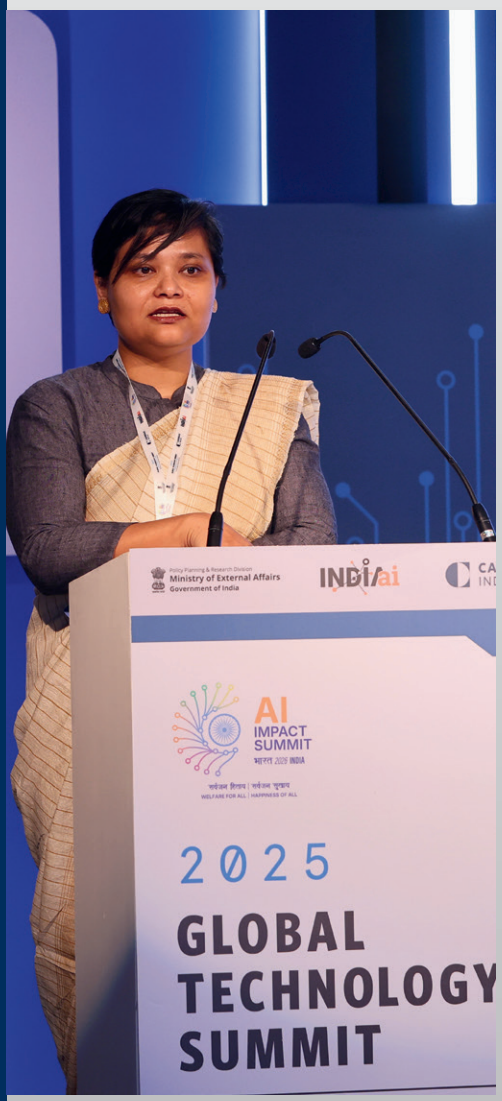

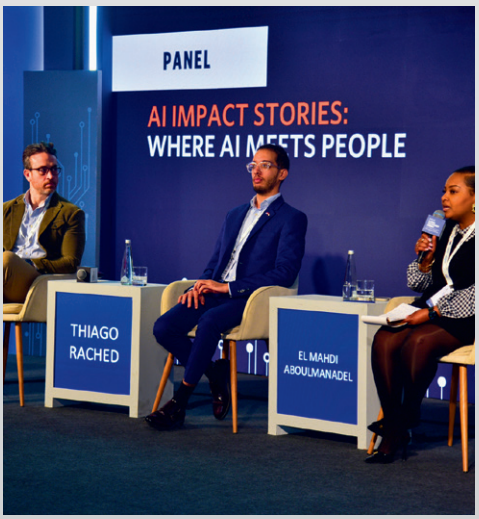

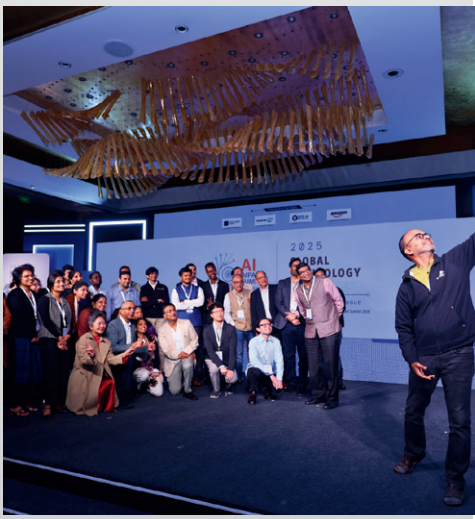

CO-HOST

Policy Planning & Research Division
Ministry of External Affairs
Government of India

PHILANTHROPY PARTNER

TATA TRUSTS

DIALOGUE PARTNER

EkStep

people + ai
An EkStep Initiative

KNOWLEDGE PARTNERS

OpenAI

Microsoft

amazon

Gates Foundation

salesforce

ANTHROP\C

LinkedIn

TATA CONSULTANCY SERVICES

CENTRAL SQUARE FOUNDATION

Meta

Policy Center for the Digital Economy

Qualcomm

THE BEST RUN SAP

HCLTech

IBM

ITU

OECD.AI Policy Observatory

AI Hub for Sustainable Development

UNDP

peak xv

MEDIA PARTNER

ANI
Asian News International

DIGITAL PARTNER

ThePrint
दिप्रिंट

OUTREACH PARTNER

asianetnews.com

CONTENTS

DISCLAIMER

Please direct inquiries to:

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

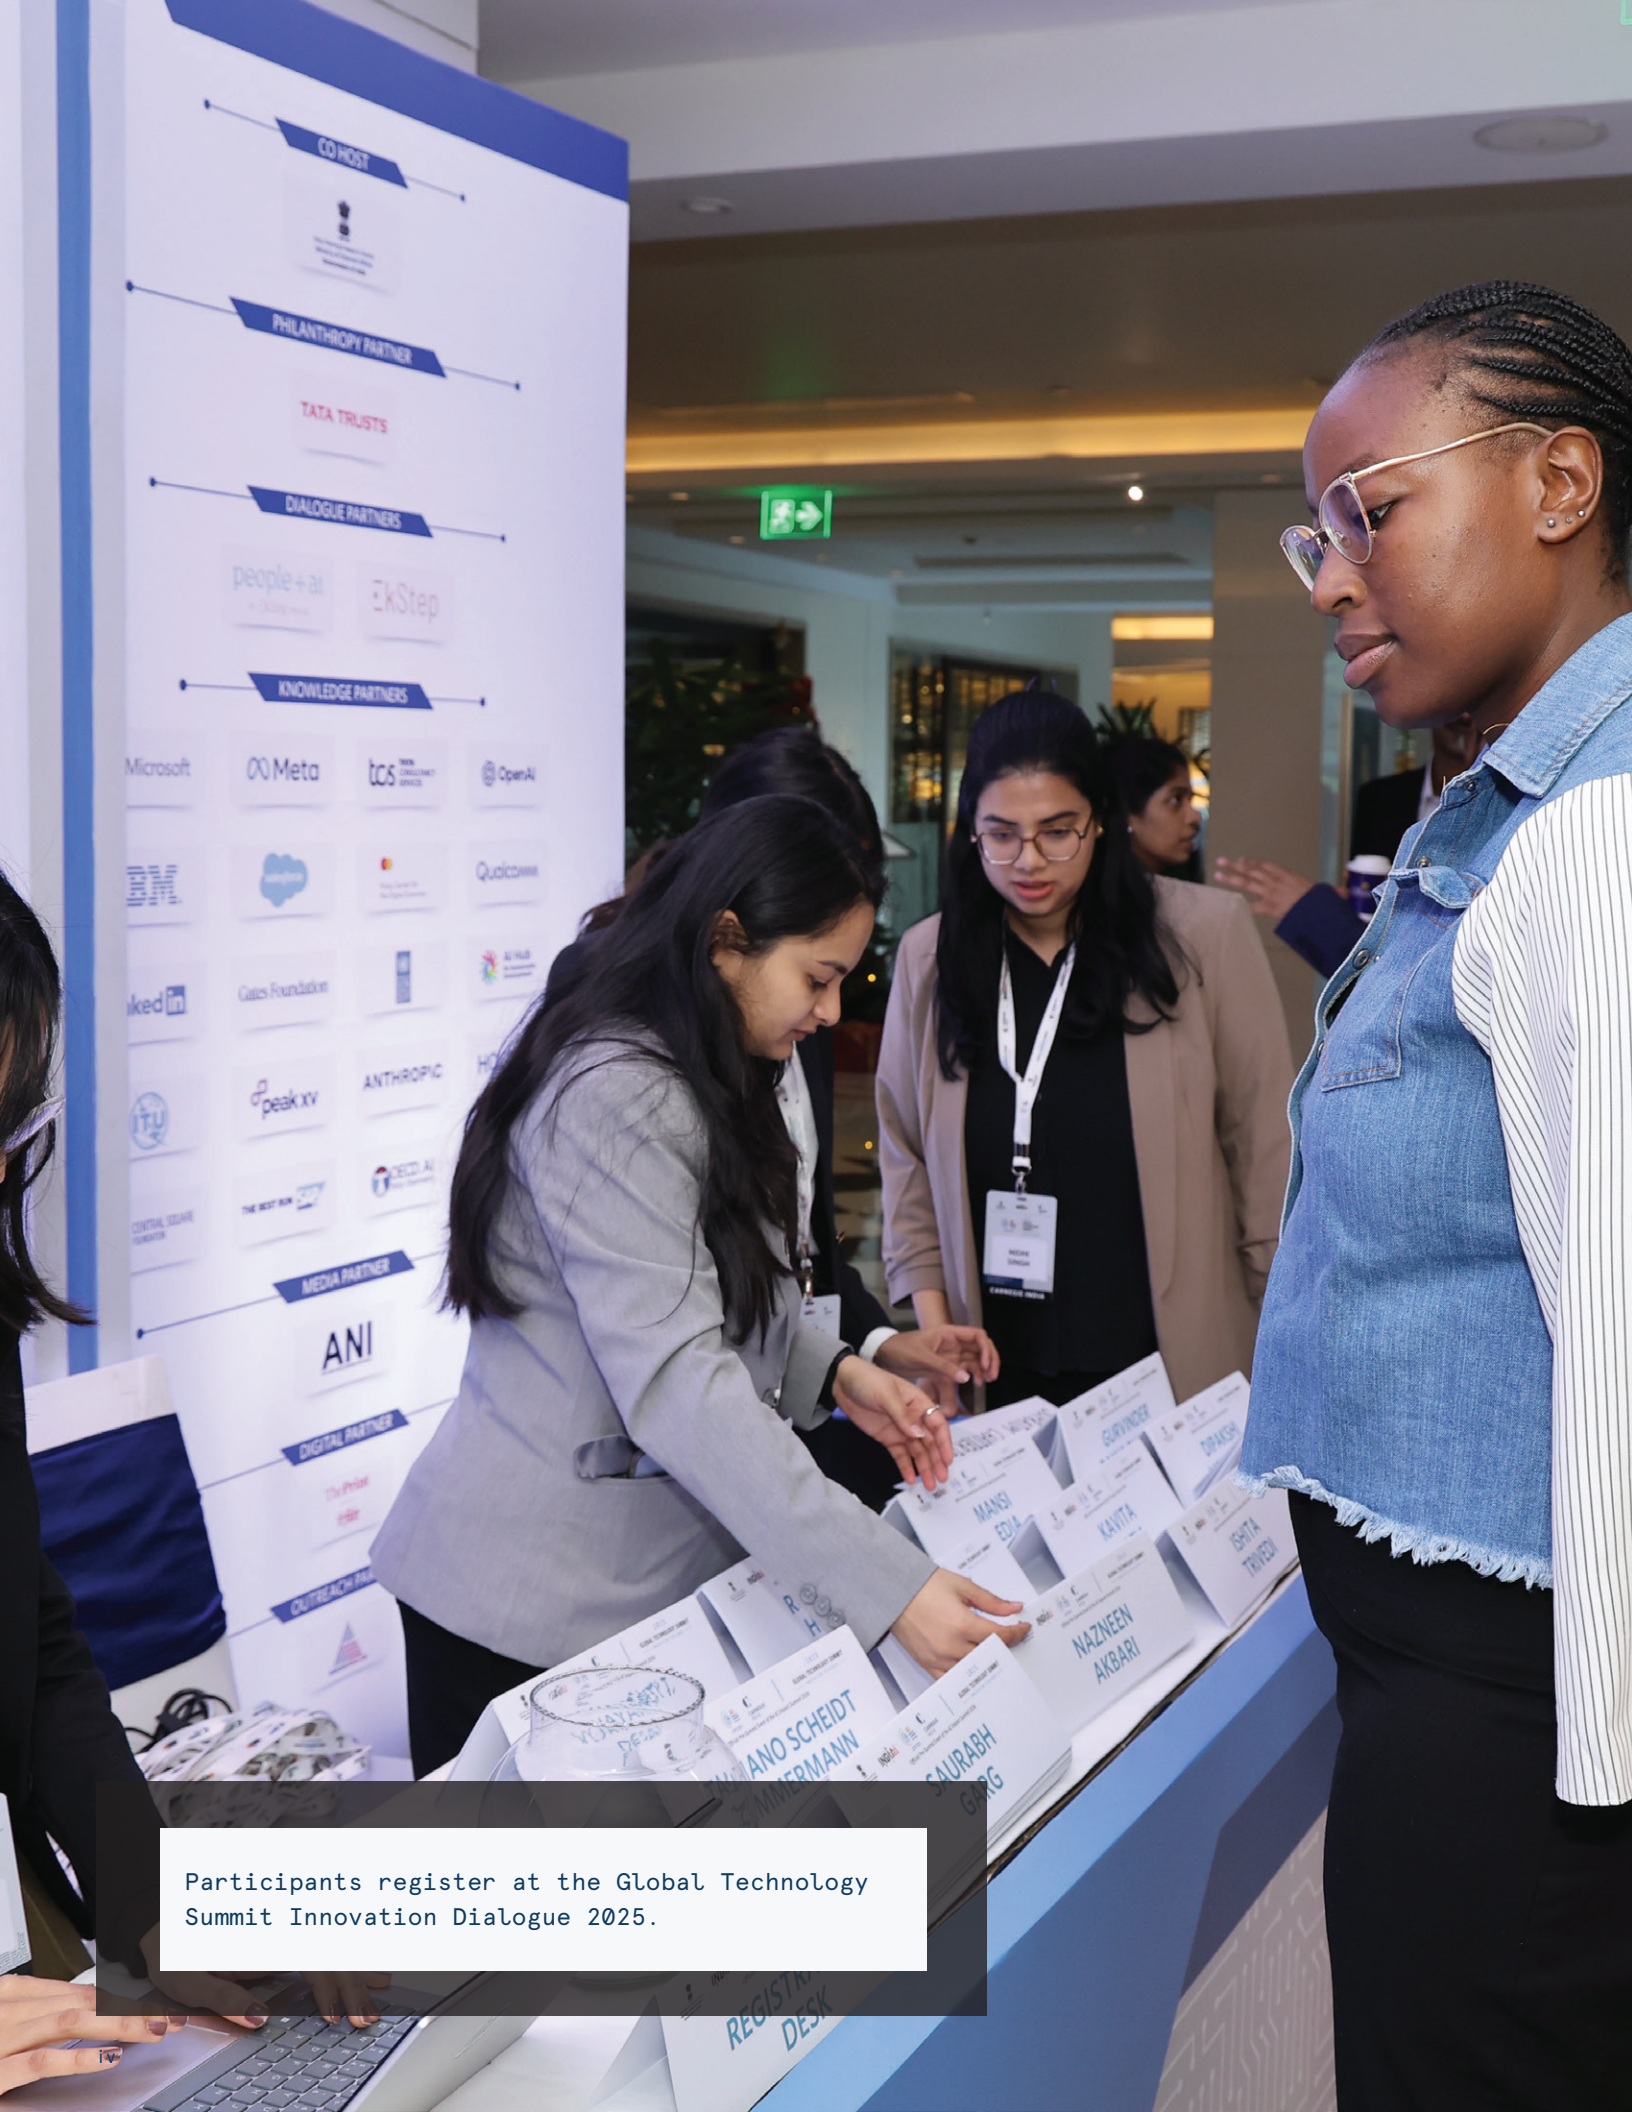

Participants register at the Global Technology Summit Innovation Dialogue 2025.

INTRODUCTION

The GTS Innovation Dialogue, held on December 10 and 11, 2025, at The Oberoi, New Delhi, was an official pre-summit event of the India AI Impact Summit 2026. Organized by Carnegie India in collaboration with the Ministry of External Affairs, Government of India, the hybrid Dialogue brought together over 200 participants across its closed-door sessions. Over 1,200 people attended the public sessions online and in-person, including senior government officials, tech experts, industry leaders, public policy professionals, entrepreneurs, civil society representatives, and scholars from India and abroad, to deliberate on AI's real-world impact across sectors such as healthcare, agriculture, and education.

The India AI Impact Summit 2026 represented the fourth iteration of global AI summits, following the AI Safety Summit (Bletchley, UK, November 2023), the AI Seoul Summit (Seoul, South Korea, May 2024, co-chaired by the UK), and the AI Action Summit (Paris, France, February 2025, co-chaired by India). Participation and influence have expanded significantly over time, from twenty-seven participating countries in 2023 to over a hundred in 2025. Structured around three core pillars, "People, Planet, and Progress," the AI Impact Summit operationalized its agenda through seven thematic working groups, with India chairing each group alongside international co-chairs. These working groups spanned AI for economic growth and social good; democratization of AI resources; science; human capital; inclusion for social empowerment; resilience, innovation, and efficiency; and safe and trusted AI. The Summit's design reflected a deeply consultative approach, incorporating engagement with over 500 organizations and global dialogue sessions in cities such as Paris, Cape Town, Berlin, Oslo, and Tokyo.

The Innovation Dialogue was timely and consequential in this context. Rather than replicating the global discourse around compute, codes of ethics, voluntary commitments, or major investment announcements, it framed India's role as host around a more consequential objective: understanding how AI meaningfully impacts people. The priority was to ensure that AI works for individuals, whether teachers, farmers, or healthcare workers, by identifying scalable use cases across regions, while simultaneously addressing regulation, governance, compute, talent, and skills. The Dialogue aimed to generate ideas, some of which could translate into outcomes for the working groups, and to contribute to a shared grammar that could inform their zero drafts, converting conversations into deliverables anchored in concrete learnings and real-world impact.

Underpinning the Dialogue was a shared recognition that the geopolitics and geoeconomics of AI's future cannot be meaningfully addressed without understanding how it is used and experienced in everyday life. There was accordingly an imperative to translate technology into practical, high-impact use cases from India and the Global South. While challenges across regions may be similar, particularly the need for technology to support inclusive growth aligned with development priorities, the solutions must remain locally shaped and context-driven. The Dialogue sought to synthesize these diverse experiences coherently, especially as global efforts move toward establishing the normative and legal frameworks that will govern emerging technologies, ensuring that India and the broader Global South play an active role in shaping these norms, rather than being treated as exceptions to them.

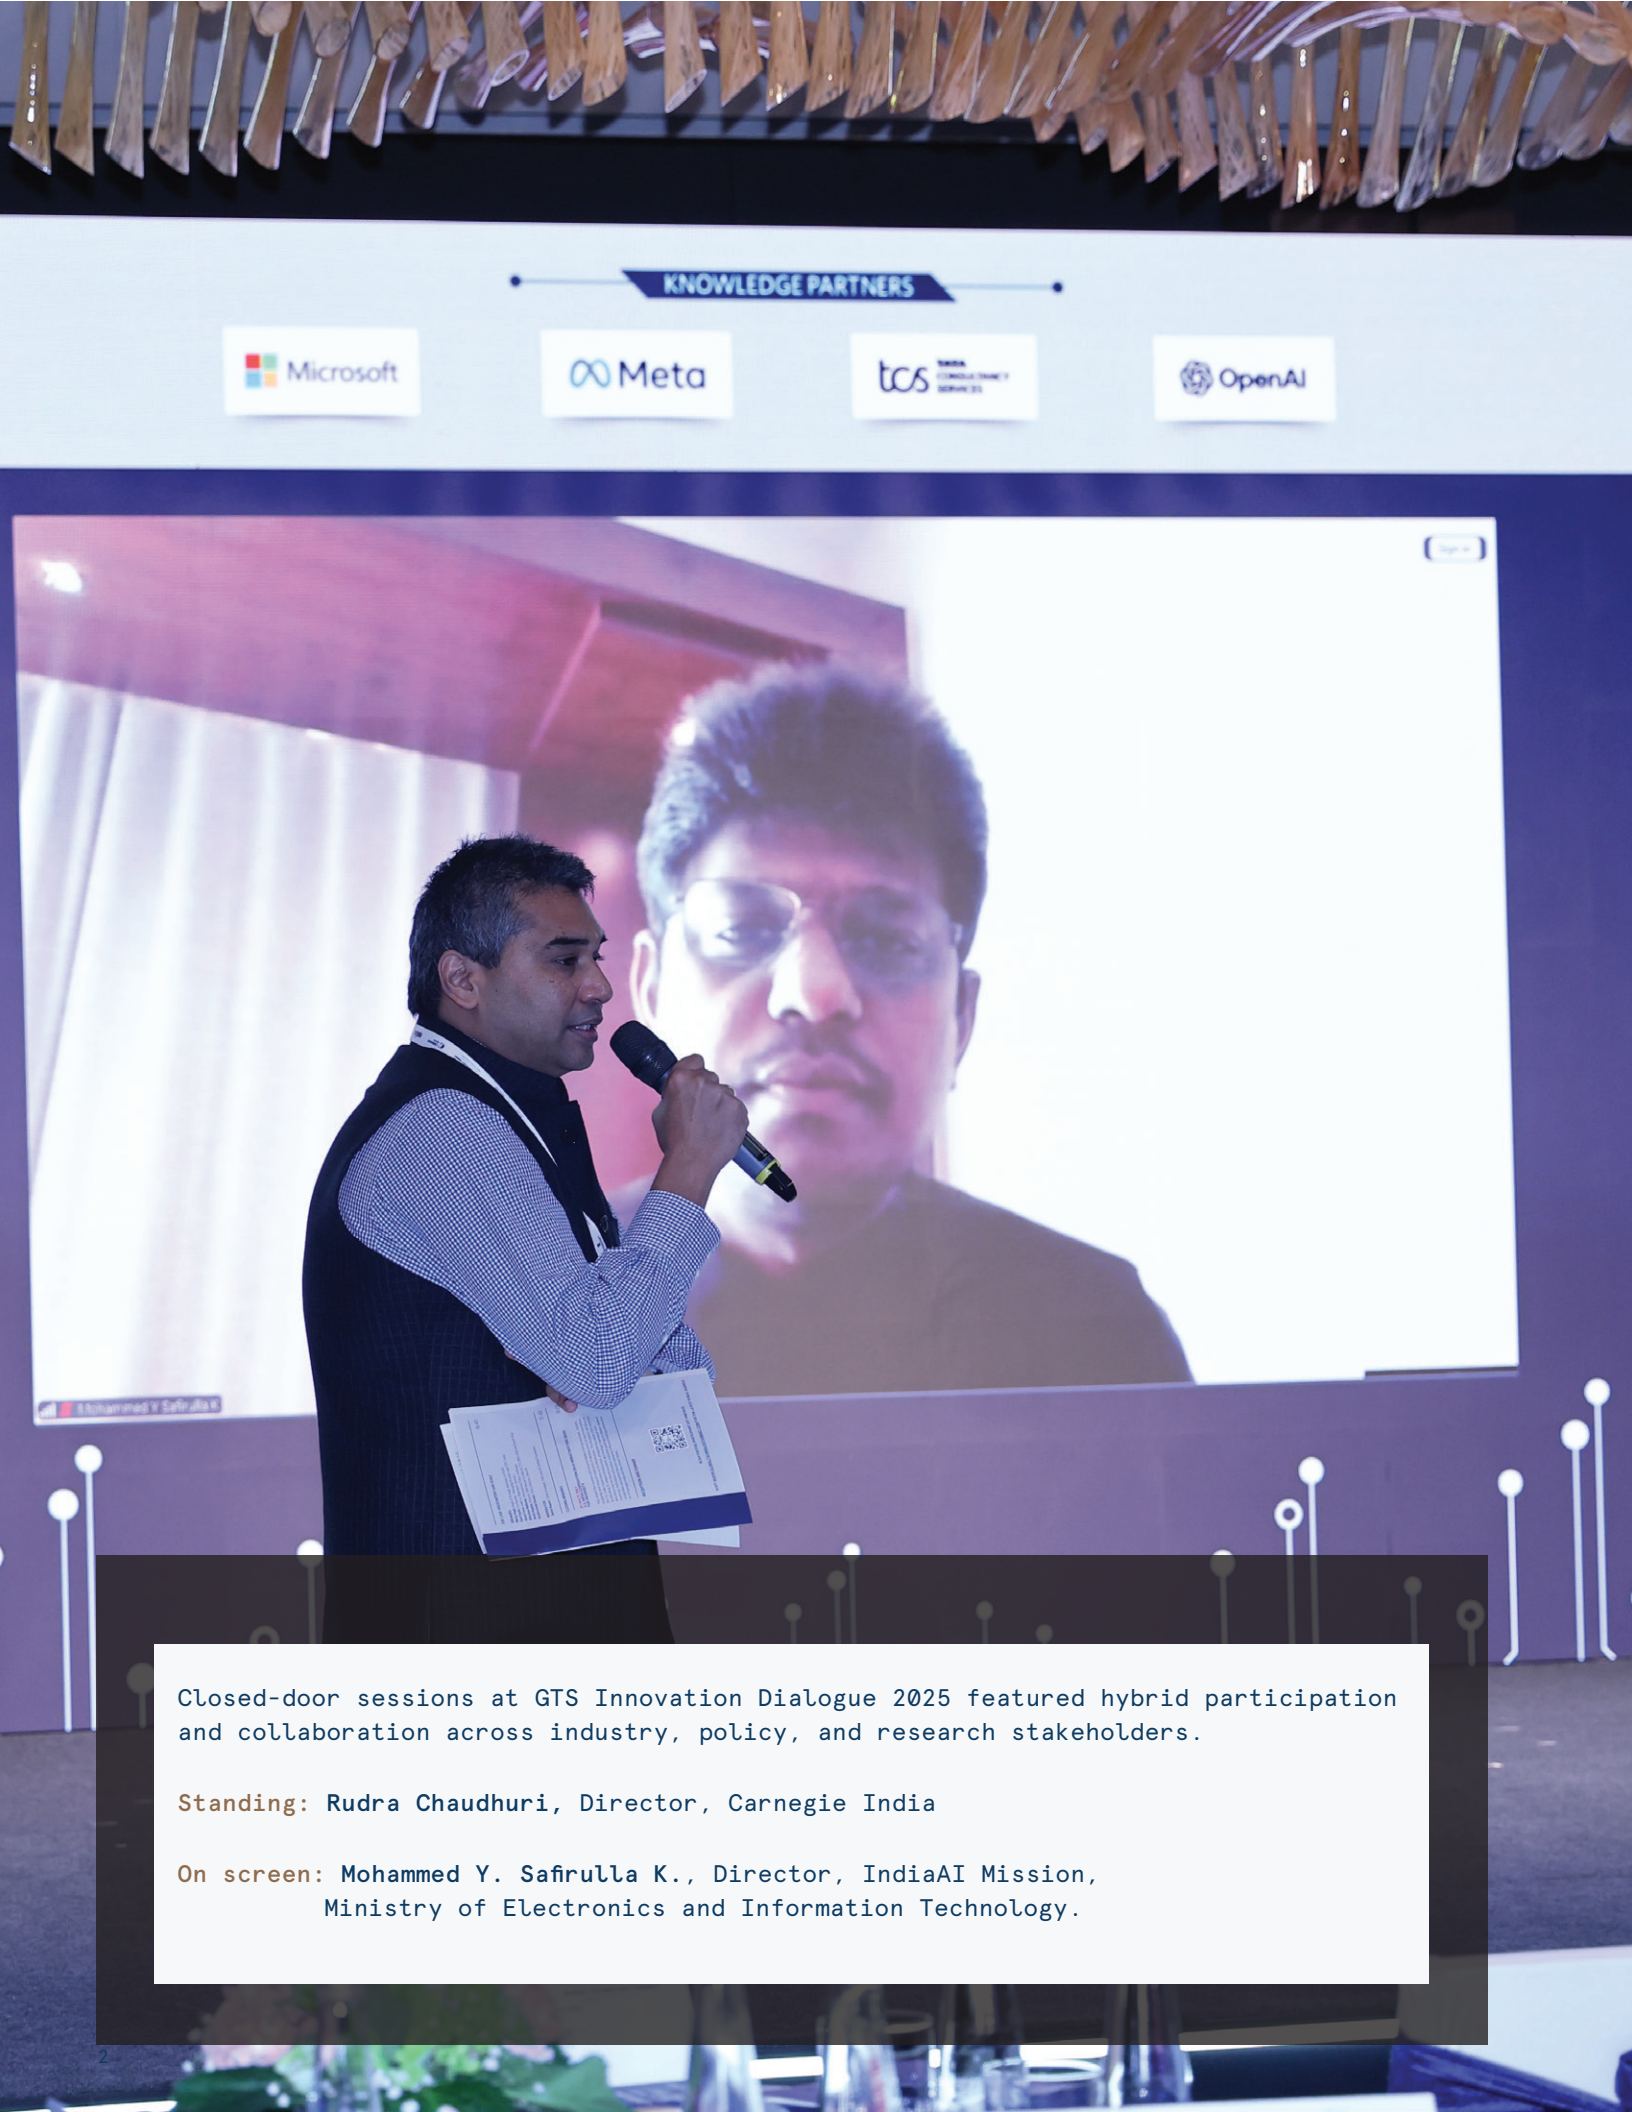

Closed-door sessions at GTS Innovation Dialogue 2025 featured hybrid participation and collaboration across industry, policy, and research stakeholders.

Standing: **Rudra Chaudhuri,** Director, Carnegie India

On screen: **Mohammed Y. Safirulla K.**, Director, IndiaAI Mission, Ministry of Electronics and Information Technology.

CLOSED-DOOR DISCUSSION HIGHLIGHTS

10 December, 2025

Challenges and Opportunities in Scaling Use Cases

The Dialogue opened with an icebreaker session that focused on key overarching questions: What did the promise of AI mean for those present? How would they measure and define impact when it came to AI use cases? The session surfaced an array of thematic areas considered fundamental to the promise of AI, such as equity, empowerment, trust, and inclusion. A parallel discussion around how impact should be assessed led to themes emerging around social metrics such as sustainability, better life, and positive transformation. Following the plenary, participants were divided into three breakout groups focused on AI in healthcare, agriculture, and education. Each group explored a similar set of guiding questions around how and why specific AI use cases were developed, how they were built in terms of sustainability and financing, and how practitioners accessed compute.

The objectives of the breakout sessions were twofold. First, to identify structural and operational barriers to AI adoption across diverse domains, looking beyond technology to systemic issues, including financing, incentive alignment, and trust. And second, to clarify what collaboration means in practice, whether through top-down platforms, network-of-networks models, or cross-border partnerships, and to explore how shared resources, including datasets, compute, and customizable models can democratize access.

Across all three breakout groups, collaboration was identified as the most pressing theme and framed as the most energy-efficient way to achieve AI goals under resource constraints. Two operational models emerged: a top-down platform that curates and coordinates, and a network-of-networks approach that grows organically and makes redundant single points of failure. Impact was reframed to go beyond adoption metrics, incorporating dignity, empowerment, and improved outcomes. While compute affordability and sovereignty remain important, they were less dominant than in other discussions on AI in the past. Cross-border opportunities, such as India-Nepal linguistic datasets, were highlighted as immediate avenues for cooperation.

Breakout Session 1: AI in Healthcare

This session explored practical applications of artificial intelligence across healthcare systems in the Global South. It examined real-world use cases spanning disease diagnosis and prevention, clinical operations, and patient care delivery, while addressing systemic barriers to implementation, including data access, workforce capacity, sustainability models, and the need for localized, context-sensitive AI solutions.

The primary objectives of this breakout session were:

1. Identify End-Uses and Scaling Journeys: The workshop sought to understand the diverse landscape of AI applications in healthcare across the Global South, from diagnostic tools addressing pathology shortages in Africa to wearable technologies enabling preventive care and supply chain solutions ensuring equitable access to medicines. Participants explored how these solutions could be scaled effectively across different regions and healthcare contexts while remaining responsive to local needs.
2. Look at Role of Datasets: The workshop identified critical barriers hindering AI implementation in healthcare sys-

Breakout Session 1: AI in Healthcare

Breakout Session 1: AI in Healthcare

tems, with particular emphasis on data access challenges, and quality and representativeness of datasets. Discussions revealed that data scarcity is often overstated; the real challenge is obtaining locally relevant, high-quality data that reflects the populations being served and understanding how much data is genuinely needed at each stage of use case development and deployment.

3. Develop Governance Frameworks that Align with Local Context: The workshop aimed to develop frameworks for responsible, localized AI deployment that account for context-specific needs, equity considerations, and long-term viability in resource-limited settings. Rather than applying global templates, the focus was on understanding how solutions must be adapted to different populations and health systems, and how validation must occur in real-world contexts.

Below are the key takeaways from this breakout session:

1. Large Technology Firms Must Work with Startups

   The most productive collaborations occur when large technology firms provide foundational models that local entrepreneurs adapt to their context. However, problems occur when technology platforms suddenly withdraw support, forcing smaller start-ups to rethink their model architecture. Furthermore, large technology firms could impose business models designed for abundance rather than resource-constrained contexts. Given this context, a lot of innovation in healthcare AI has come from start-ups that operate under constraints of frugality.

2. More Incentives Needed for Last-Mile Health Solutions

   Tropical diseases, mental health in the Global South, early childhood intervention for preventable disabilities, and AI solutions for populations with disabilities all represent areas where commercial incentives are weak, but health impact is enormous. Funding mechanisms must actively create incentives to develop solutions through preferential procurement, dedicated data collection funding, and recognition systems that reward innovations to improve underfunded and under-recognised conditions.

3. Sustainable Financing Needs a Long-Term Roadmap

   Current funding patterns (donors financing startups or government agencies, without guaranteeing government adoption) create the "last-mile problem" where promising tools cannot scale because health systems lack budget and integration pathways. Solutions include restructuring funding around "data as a service" models, requiring government co-financing commitments upfront, and moving procurement beyond software licenses to holistic risk-assessment frameworks that account for workflow integration, liability, and ongoing monitoring.

In the case of AI deployments in healthcare, there was recognition that compute needs vary drastically across the lifecycle, from pilot to deployment, and require careful planning. Consortia were identified as critical for validation, benchmarking, and supporting start-ups at different growth stages. Unlocking sectoral data and aligning incentives were seen as key to accelerating adoption, while consortium support must evolve to meet the distinct needs of early-stage and growth-stage innovators.

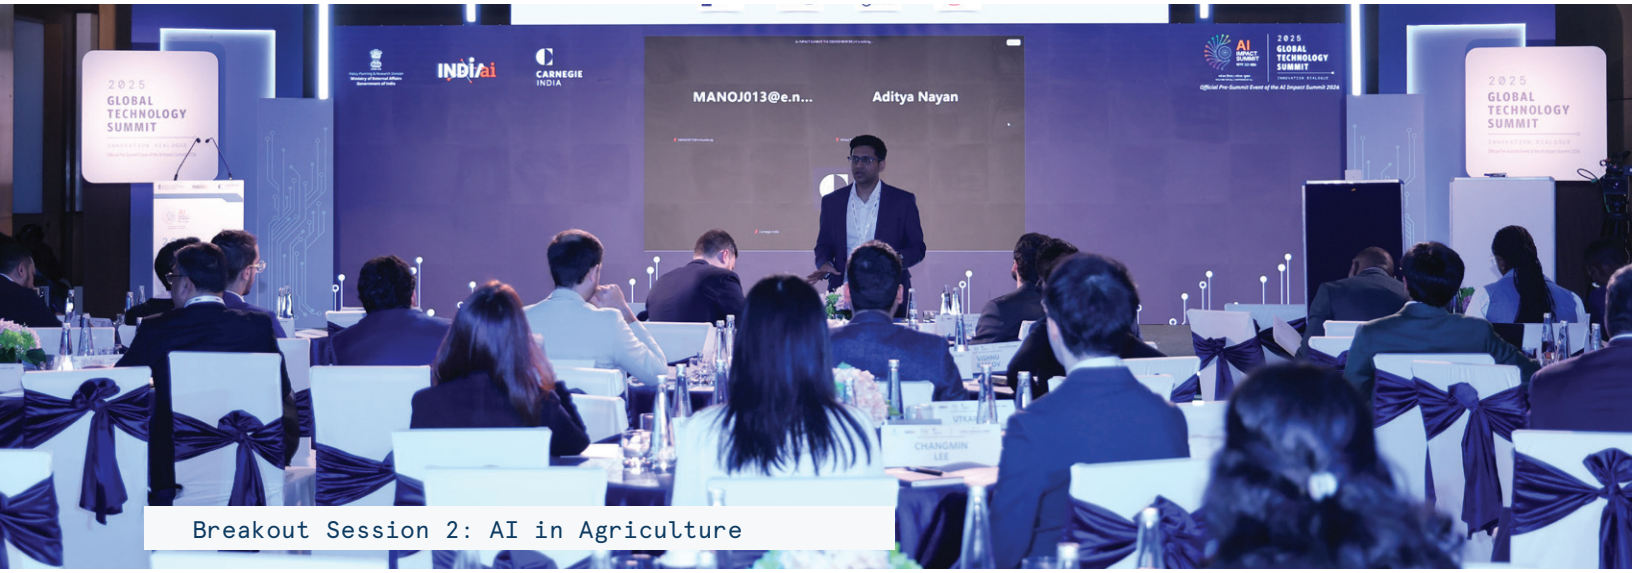

Breakout Session 2: AI in Agriculture

Breakout Session 2: AI in Agriculture

The breakout session on AI in Agriculture brought together use case developers, hyperscalers, think tanks, and ecosystem stakeholders to examine practical challenges and opportunities in scaling agricultural AI solutions. Participants identified common challenges that transcended geographical boundaries with the aim to develop common and workable solutions. Participants shared and reflected on real-world experiences of building AI solutions across diverse contexts.

The primary objectives of this breakout session were:

1. Document Implementation Journeys: Capture first-hand accounts of how agricultural AI use cases were conceptualized, developed, and deployed, including the institutional support mechanisms, partnerships, and problem-solving approaches that enabled their creation.
2. Map Resource Access Pathways: Understand how practitioners navigated critical resource constraints, specifically dataset acquisition and training, access to AI infrastructure, and financing, to identify replicable models and highlight persistent barriers requiring systemic intervention.
3. Identify Scaling Bottlenecks and Solutions: Through cross-stakeholder dialogue between developers, infrastructure providers, and policy actors, pinpoint the specific technical, financial, and operational challenges that prevent promising pilots from achieving scale and sustainability.
4. Build Collaborative Networks: Foster connections among hyperscalers, use case implementers, think tanks, and support organizations to facilitate knowledge sharing, resource pooling, and coordinated approaches that can accelerate agricultural AI adoption across India's diverse farming landscape.

Below are the key takeaways from this breakout session:

1. Hyperlocalization is Non-negotiable

   AI solutions must be tailored to specific geographies, weather patterns, crop cycles, and pest profiles. This requires investment in local language datasets, voice technologies, and context-aware models that account for agricultural diversity while managing the associated demands on data and compute.

2. Hybrid Compute Strategy Essential

   Each AI application has varying compute requirements. Massive amounts of compute is not necessary for all applications. A balanced approach combining high-performance computing for model development with edge-based inference could help address connectivity limitations, sovereignty concerns, and cost constraints while enabling fine-tuning for localized conditions.

3. Financing Must Center Around Farmers

   Sustainable scaling requires blended finance models that combine public-private partnerships, outcome-based payments, and crop-cycle credit mechanisms. Loan structures could include provisions for compute costs, while monetization should target private sector value capture rather than farmer payments. Crop insurance integration and patient capital are critical for risk mitigation.

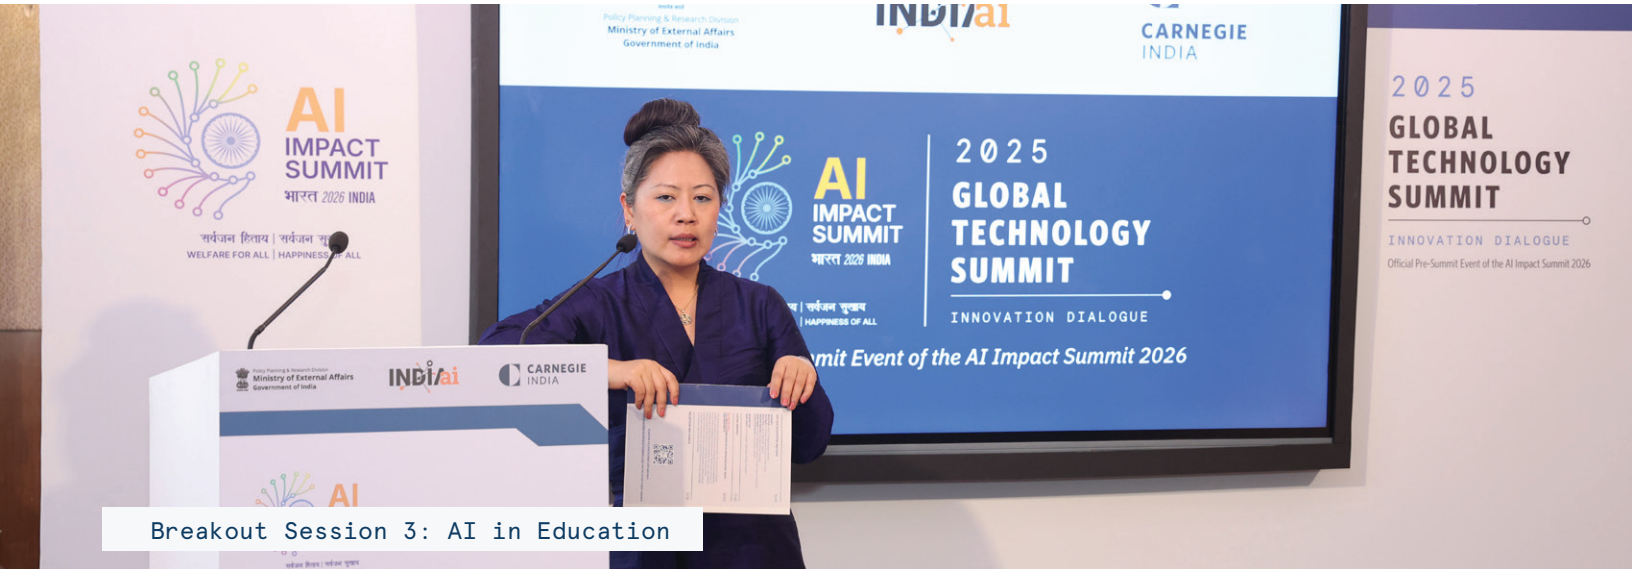

Breakout Session 3: AI in Education

4. Development of Reliable Data Infrastructure is a Pre-requisite

   Effective use of AI in agriculture hinges on building reliable, transparent, and farmer-centric data ecosystems. Data must be closely linked to market signals where the context of the dataset is locally and commercially relevant and can support timely, revenue-informed decision-making, while also enabling continuous knowledge transfer through feedback loops that incorporate farmer experience and local knowledge. Ensuring data reliability, clear accountability among data collectors, and verifiable records is critical for trust in AI systems.

5. Trust Building is Foundational

   Wide AI adoption depends on explainability, institutional support, and capacity building that helps farmers understand and evaluate AI recommendations against traditional knowledge. Government policies must address any risk aversion here among farmers by creating accountability mechanisms that validate AI performance over multiple crop cycles.

The session further underscored the importance of hyper-local solutions, sometimes within a 1 to 10 km radius, and the need for time-sensitive deployment aligned with seasonal cycles. Farmers must remain central to AI solutions, though the wider ecosystem cannot be ignored. Compute costs were flagged as a major challenge, with examples showing the high expense of serving large populations. Participants called for readiness frameworks, explainable models, and financing mechanisms that allow solutions to start small while remaining accessible and trustworthy.

Breakout Session 3: AI in Education

This session brought together practitioners, researchers, funders, technologists, and policymakers to explore how AI can meaningfully support education systems, particularly in diverse, resource-constrained, and multilingual contexts. Through introductions and discussion, participants surfaced real-world challenges, shared live use cases, and debated how AI can scale responsibly while remaining grounded in the needs of learners, teachers, and community as a whole.

The primary objectives of this breakout session were:

1. Augmentation of Institutional Capacity: To examine how AI can address concrete educational pain points, especially those faced by teachers and learners in overcrowded classrooms, under-resourced schools, and multilingual environments. It emphasized problem-first thinking by focusing on literacy, numeracy, feedback, remediation, and teacher workload.

2. To Understand the Constituents of Scale: To understand what "scale" meant in education, particularly in the Global South. The breakout session here was keen to explore how infrastructure, cost, language, dialects, cultural context, and access to devices collectively shape whether AI solutions can realistically reach population scale.

3. To Explore Responsible Use of AI in Education: To surface shared principles for responsible and inclusive AI in education. Participants reflected on issues such as hallucinations, over-reliance on AI, loss of critical thinking, benchmarking bias, and the need to preserve "cognitive sovereignty" by ensuring AI supports learning how to think, not replacing human judgment or agency.

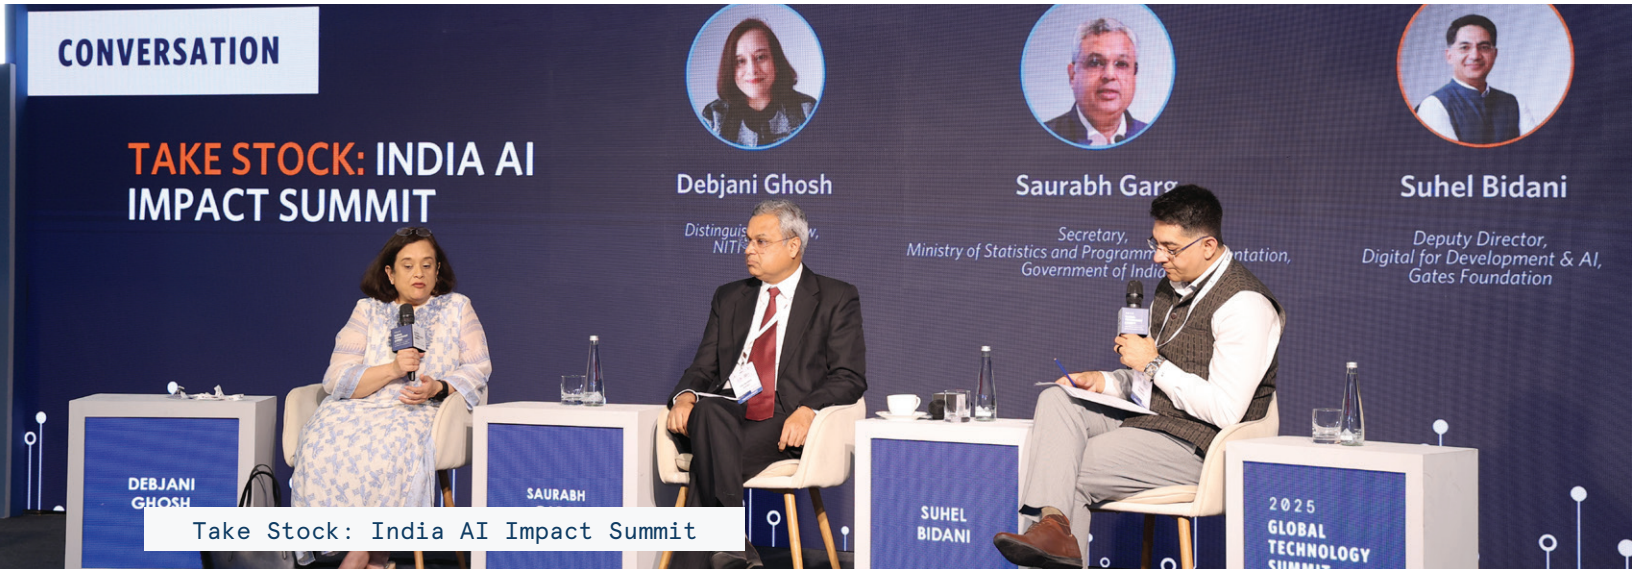

Take Stock: India AI Impact Summit

Below are the key takeaways from this breakout session:

1. Adopt a Problem-Centric Approach

   Focus on solving specific, high-priority pain points of teachers, such as grading burdens and remediation, by designing AI models that work as a "co-pilot" that augments teacher capacity, rather than replacing them altogether.

2. Build Localized Infrastructure and Networks

   Establish collaborative networks for affordable compute, leverage cross-sectoral expertise, and aim to build capacity to ensure the Global South has the local infrastructure and technical literacy to lead its own AI development.

3. Prioritize Contextual R&D and Evaluation

   Instead of adapting imported Global North benchmarks, there was emphasis on investing in Global South benchmarks led by research that prioritizes local languages, dialects, and cultural nuances through community-driven data and evaluation.

4. Innovative Funding and Business Models Required

   There needs to be a shift towards philanthropic-first and B2G (Business to Government) models that favor local impact over hyper scaling, using affordable micro-transaction or "bundle" pricing to make AI viable in price-sensitive markets. Scaling was not always regarded as a priority in AI applications for education, in comparison to the efficiency and effectiveness it offered to students, parents, and teachers.

5. Ensure Ethical and Cognitive Sovereignty

   The breakout session advocated for proactively addressing technical risks such as hallucinations and data privacy while ensuring the availability of AI tools that supported students with "learning how to learn" rather than creating a dependency that led to the erosion of critical thinking. There was an emphasis on AI supporting "thought processes" and reducing the burden on teachers in the education sector, rather than AI replacing "thought processes" and teachers.

A key finding of the session was that standalone solutions will not work in education. Incentive alignment and continuous collaboration are essential to sustain AI use cases. Solutions must be modular and layered, tailored to local contexts, and designed to unlock systemic barriers rather than treat AI as standalone tools. Co-creation and iterative problem-solving were seen as essential, even when serving smaller populations. Networks were highlighted as recurring enablers, accelerating adoption and scaling solutions where appropriate.

Take Stock: India AI Impact Summit

This conversation brought together co-chairs from two working groups under the AI Impact Summit: Ms. Debjani Ghosh, distinguished fellow at the NITI Aayog who leads the work on "AI for Economic Growth and Social Good;" and Dr. Saurabh Garg, secretary of the Ministry of Statistics and Programme Implementation (MoSPI), Government of India, who is leading the Working Group on "Democratizing AI Resources."

The conversation focused on moving beyond traditional narratives around AI's promise to concrete, scalable impact, particularly for economic development, social inclusion, and national resilience. Speakers emphasized collaboration across

countries, shared infrastructure and knowledge, and the need to institutionalize outcomes beyond the Summit itself.

Below were the primary objectives of the conversation.

1. To Take Stock: To understand the progress and status of working groups on "Economic Growth and Social Good," and on "Democratizing AI."
2. To Understand What End-Users Need: To focus on operating toolkits and use cases, particularly for the Global South, to enable cross-border collaboration and knowledge-sharing.
3. To Manage Expectations: To understand the support expected from non-profit organizations, academia, developmental entities, and countries.
4. To Ensure Continuity of Outcomes: To help institutionalize post-summit mechanisms so that tools, best practices, and partnerships endure beyond the AI Summit.

Below are the key takeaways from this conversation:

1. Impact at Scale as the Core Objective

   There is strong convergence on the need to re-position AI from a narrative of promise to a framework of delivery. AI is increasingly framed as an enabler of economic growth, social development, technological sovereignty, and national resilience, not an end. There was an emphasis on moving beyond pilots and isolated experiments toward measurable outcomes at scale, with impact creation established as the primary metric of success.

2. Global Collaboration to Accelerate Replication and Scale

   Achieving impact at scale requires collective action. No country can afford to develop solutions in isolation or repeatedly solve the same problems. Shared platforms, such as an emerging "Impact Commons," aim to systematize learnings, best practices, and toolkits so that countries can adapt proven approaches to local contexts. This collaborative model is particularly critical for emerging markets and developing economies seeking to accelerate development outcomes efficiently.

3. Democratizing AI Through Shared Resources and Institutional Design

   Democratization of AI extends beyond access to technology. While compute, data, models, and talent remain foundational, equal emphasis is placed on connectivity, openness, community engagement, and governance. Proposed mechanisms such as a "Charter for AI Diffusion" reflect growing consensus around shared infrastructure, open innovation, blended financing, and respect for data sovereignty. Crucially, democratization must be institutionalized beyond the summit cycle to ensure long-term affordability, continuity, and accessibility, especially for the Global South.

4. From Use Cases to Transferable Impact Models

   There is a deliberate shift away from showcasing isolated use cases, toward documenting comprehensive impact stories. These emphasize the problem being addressed, the scale and sustainability of outcomes, design principles used, challenges encountered, and pathways for replication. By foregrounding impact rather than technology alone, such models become actionable across diverse national and sectoral contexts.

5. Governance Challenges: Trust, Jobs, and Concentration of Power

   Several structural challenges emerged as central to AI governance. Trusted AI depends on technical safeguards for safety and security and on protocol-based credibility for cross-border data sharing and collaboration. At the same time, AI-driven job displacement, particularly in mass-employment, entry-level sectors, is acknowledged as inevitable and insufficiently addressed, requiring urgent policy attention. Concerns around the concentration of AI capabilities among large technology firms further underscore the need for shared resources, inclusive datasets, grassroots innovation, and people-centric approaches to ensure AI remains equitable and globally distributed.

AI Use Case (Dis)content and Scale

This discussion brought together infrastructure providers, model developers, and AI companies to examine the practical challenges of scaling AI use cases. The discussion focused on how India has evolved from compute scarcity to facing new bottlenecks in the "missing middle layer" of deployment infrastructure, while exploring collaborative approaches to build complete technology stacks, address multilingual performance gaps, and create sustainable business models that enable widespread AI adoption across agriculture, healthcare, and education sectors.

The primary objectives of the panel discussion were:

1. Identify Adoption Bottlenecks: To unpack why AI adoption remained difficult even when compute, models, and

AI Use Case (Dis)content and Scale

ideas are available, and to identify the key bottlenecks that prevented pilots from translating into scalable, real-world deployments, particularly the missing "middle layer" of execution, orchestration, data integration, and institutional capability.

2. Explore Collaborative Ecosystem Solutions: To examine how collaborative ecosystems can address these gaps in the Indian and Global South context. The panel discussion was also a way to explore sustainable compute economics, multilingual and cultural alignment, responsible model development, and the role of government-led benchmarks.

Below are the key takeaways from the panel discussion:

1. The Missing Middle Layer Requires Urgent Attention

While India has, for now, made an effort to address compute scarcity and generated abundant use case ideas, the critical gap lies in execution infrastructure. Organizations lack middleware for deployment at scale, orchestration tools to coordinate system components, fine-tuning expertise for specific domains, practical Retrieval-Augmented Generation (RAG), implementation knowledge, and solutions for data curation across dispersed databases. Building communities of practice around these scaling architectures and creating reusable components for common use cases represents the next frontier for ecosystem development.

2. Sustainable Infrastructure Economics Must Balance Multiple Stakeholders

The IndiaAI mission successfully catalyzed ecosystem growth by guaranteeing demand for GPU providers and making compute accessible to start-ups, enabling the development of India-specific LLMs that would otherwise be capital-prohibitive. However, extremely low pricing from competitive bidding (₹65 per GPU hour) and short one-year contract terms threaten provider viability and long-term ecosystem health. Future procurement must balance affordability with sustainable pricing for researchers and start-ups to maintain infrastructure investment confidence.

3. Government-Led Benchmarking Can Drive Language and Domain Progress

Companies cannot be solely responsible for preserving low-resource languages or achieving domain-specific performance as commercial incentives naturally misalign with these goals. Governments should create and maintain high-quality, domain-specific benchmarks (agriculture, healthcare, education, legal systems) covering local languages and cultural contexts, then publish comparative results across all domestic and foreign models. This creates competitive pressure that incentivizes companies to improve performance without relying on corporate goodwill alone.

4. Multilingual Tokenization Inequality Undermines Democratization

Fundamental technical inequalities exist in how major language models handle non-Western languages. In tokenizers containing 1, 30,000 tokens, only 6,000–10,000 may be allocated for languages like Hindi or Marathi, significantly increasing the cost of processing these languages and directly contradicting democratization goals. Models like Apertus demonstrate best practices by ensuring fair tokenization across all supported languages, including low-resource languages with only 60,000 speakers, while encoding the complete cultural context that makes downstream applications more cost-effective and contextually appropriate.

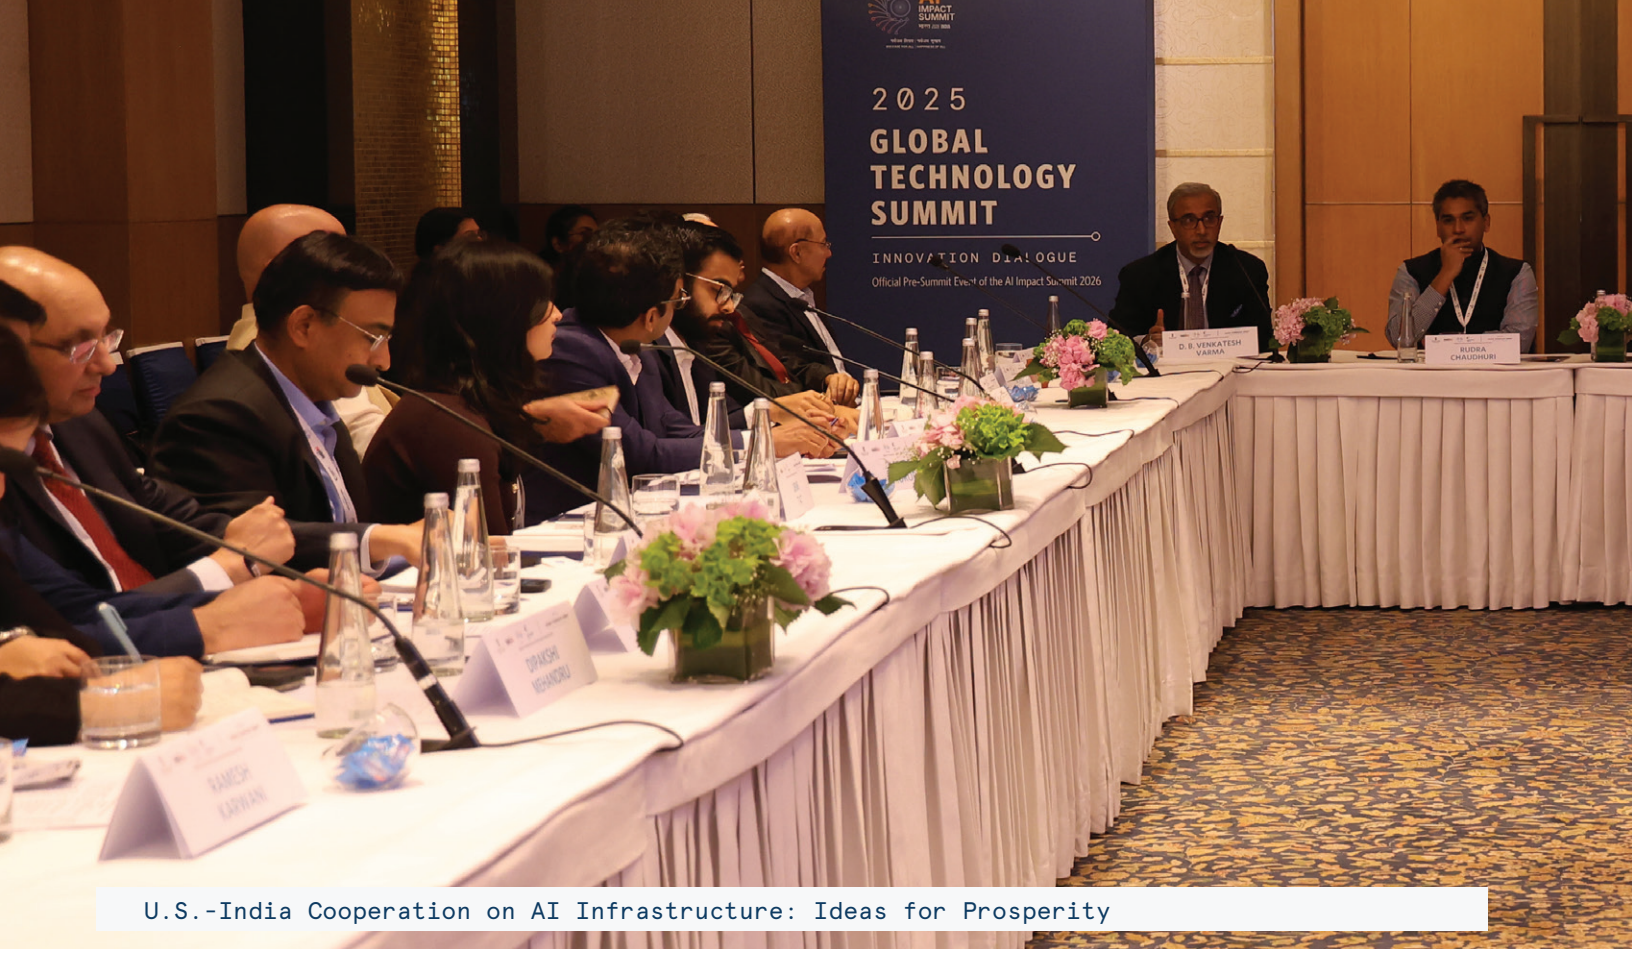

5. Real-World Deployment Accelerates Learning Beyond Laboratory Testing

   To discover what works in practice, there is no substitute for putting AI systems into the hands of actual end users-farmers, healthcare workers, students, and citizens. The history of ChatGPT illustrates this: the underlying model existed over a year before consumer release, but real-world usage drove rapid improvements that made it genuinely useful. India holds particular advantages for this approach, given its higher AI-native adoption compared to its Western counterparts. The recommended strategy begins with small pilots in specific sectors, structures them to manage risk while creating space to learn from failures, then rapidly iterates based on genuine user feedback before scaling successful patterns.

U.S.-India Cooperation on AI Infrastructure: Ideas for Prosperity

The closed-door discussion on "U.S.-India Cooperation on AI Infrastructure: Ideas for Prosperity" was attended by senior government officials, industry executives, representatives from embassies, and non-profit organizations, who have been closely involved in shaping the U.S.-India cooperation on AI. During the discussion, issues faced in the U.S.-India partnership and opportunities present concerning technology cooperation, AI infrastructure development, regulation of AI, and supply chain resilience were addressed.

The primary objectives of the closed-door discussion were:

1. Identify Bottlenecks and Solutions: The purpose was to review ongoing U.S.-India initiatives (public and private), track bilateral investments and partnerships, and identify hurdles with regard to high-level approvals, regulatory clarity, or policy alignment to unlock further progress. The roundtable also aimed to assess whether various stakeholders can engage to cooperate further on AI infrastructure development, knowledge sharing, and resource pooling.

2. Explore Future Areas of Bilateral Collaboration and Architecture Building: A core idea was to discuss next steps for cooperation across compute, AI safety institutes, supply chains, civil nuclear energy, AI infrastructure, regulatory sync, collaboration between governments, and the role of the private sector.

3. Define India's Distinct AI Narrative and Global Role: There was a need to understand and articulate India's emerging approach to AI, which was centered on democratization, societal impact, and developmental use cases. At the same time, there was a need to identify the convergences and divergences on AI between the U.S. and India, enabling public, private, and government cooperation for the India AI Summit.

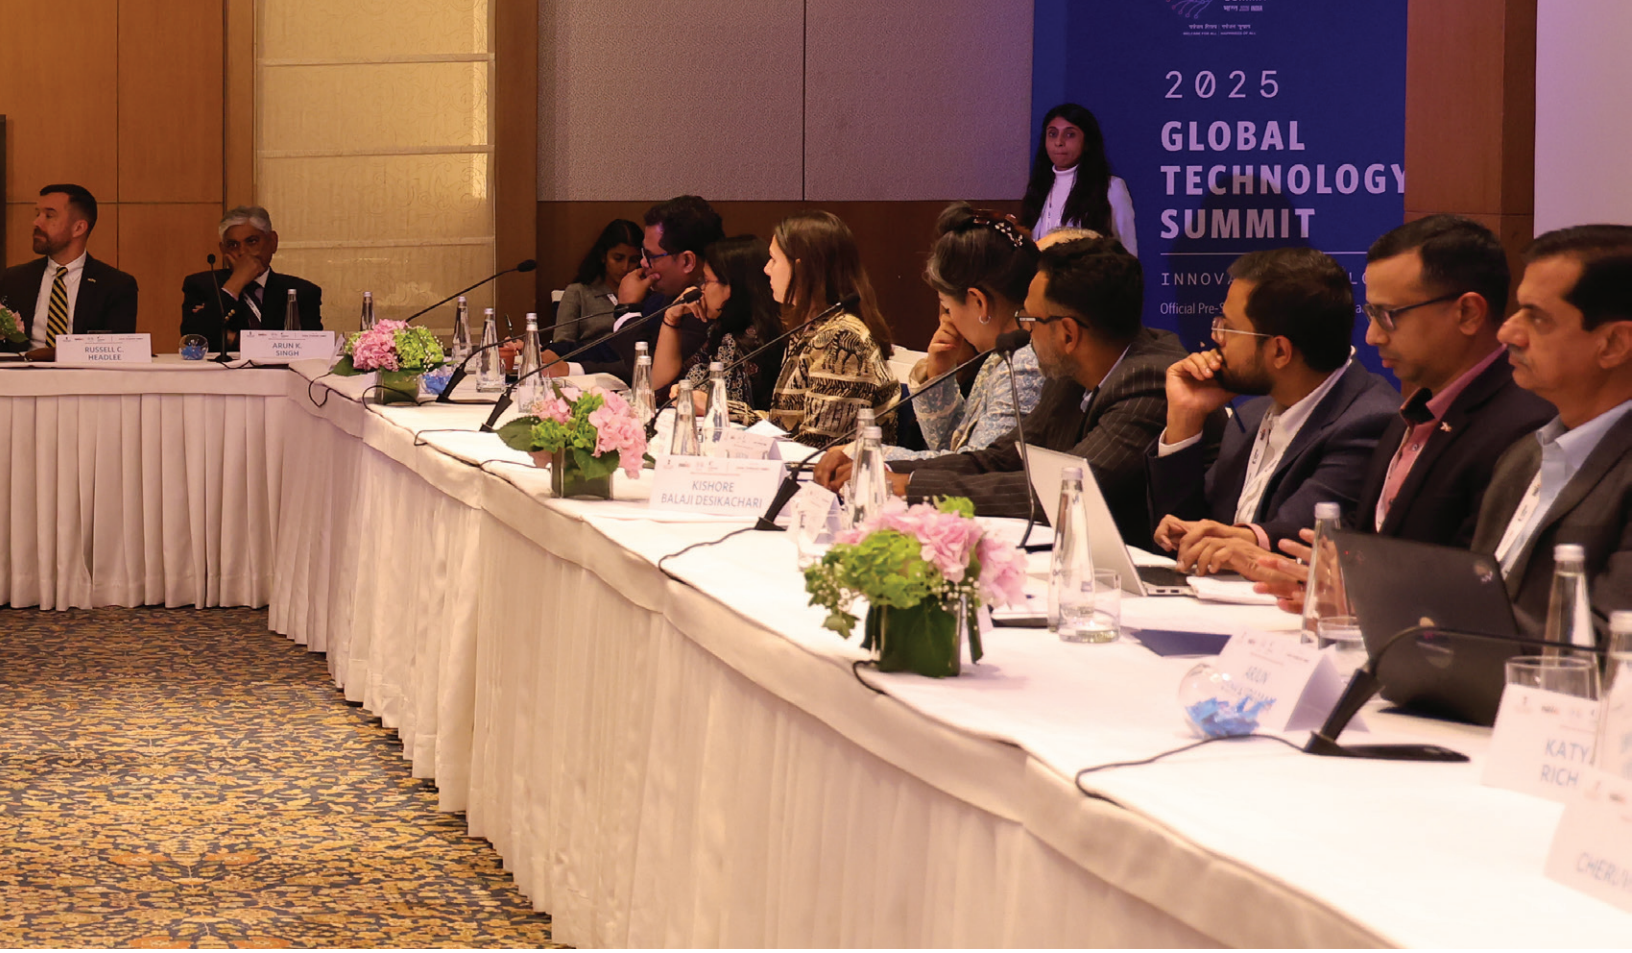

Below are the key takeaways from the closed-door discussion:

1. Operational Progress vs. High-level Sign-off

   U.S.-India cooperation on AI and advanced technology is progressing actively at the operational level, with various ministries, agencies, and companies driving their initiatives forward. However, without timely high-level sign-offs, there is a risk that operational progress could slow or stall. Both sides recognize that the strategic fundamentals of the bilateral relationship remain strong while acknowledging that lingering political hurdles need clearing for momentum to fully translate into outcomes.

2. U.S. Private Sector Investments as the Key Pillar

   U.S. technology companies are making substantial long-term investments in India across R&D, data centers, skill development, AI products, AI safety, and job creation. These investments form a vital pillar of the bilateral relationship and need to be formally acknowledged as a strategic asset, reinforcing trust, deepening collaboration, and demonstrating the tangible benefits of U.S.-India cooperation.

3. India's AI Narrative: Democratization, Relevance, and Impact

   Global AI discourse has evolved from safety and national security to governance and regulation over the course of various successive major international AI summits. India intends to add a distinct emphasis on democratization of access, ensuring AI benefits are broadly diffused, socially meaningful, prioritize real-world impact, and aligned with market relevance. By focusing on public goods innovation and impact at the population scale, India seeks to shape a new normative benchmark within emerging AI governance architectures.

4. The Need for Clear Grammar in India's AI Positioning

   A critical requirement ahead is the formulation of a clear, coherent, and consistent articulation of India's AI strategy and value proposition. Developing the right "grammar," including conceptual clarity around diffusion, social impact, persona, values, and public-good framing, will strengthen India's narrative globally and make it easier for partners and markets to engage.

5. Nuclear Energy as an Enabler for AI Infrastructure

   India's growing AI compute and data center ecosystem is driving major energy demands, making reliable baseload power a strategic priority. As nuclear energy emerges as a key enabler, India is removing civil nuclear energy investment obstacles and opening space for private-sector participation, potentially creating new opportunities

for U.S.-India cooperation in AI infrastructure and energy security.

11 December, 2025

Use Case Framework for AI

This closed-door session sought to introduce a framework that could be used to structure and augment AI use cases intended for beneficial deployment across sectors such as health, education, agriculture, and governance. The session focused on how government officials, funders, and ecosystem actors could use the framework to examine the underlying logic of an AI use case. This includes its objectives, dependencies, risks, and accountability pathways, before it was institutionalized by respective ministries, courts, welfare agencies, or service-delivery programs. Using examples from agriculture, migration, and the justice system, the session facilitators explored how the framework could help assess individual AI applications before institutions decided whether to adopt them.

The primary objectives of the closed-door session were:

1. Present a Framework for Distinguishing Experimentation from Investable AI Use Cases: To introduce a structured way of assessing when an AI application moves beyond experimentation and merits public or private investment, based on problem clarity, user relevance, and system integration.
2. Use the Framework to Examine How Governments Prioritize AI Interventions: To apply the framework to understand how public institutions evaluate AI use cases under constraints of resources, capacity, and risk, and why some interventions advance while others remain exploratory.
3. Surface Governance Requirements Through the Framework Lens: To use the framework to identify governance conditions required for sustained deployment, including procurement pathways, data stewardship, coordination across agencies, and accountability mechanisms.

Below are the key takeaways from this session:

1. Government Policy Should Focus on Specific Outcomes

   Speakers stressed that governments should adopt interventions tied to specific policy objectives, not "AI" in the abstract. Effective use cases require early clarity on the problem being addressed, the intended beneficiaries, and the decision or service being supported, before technical design begins.

2. A Single-Forum is Needed to Coordinate Multiple AI Initiatives

   Many AI initiatives stall because responsibility for data, funding, implementation, and oversight are split across multiple agencies. In the absence of clear institutional ownership, even technically robust use cases struggle to progress.

3. AI Pilots Should Test More than Technical Parameters

   While pilots often show that an AI model can work, they rarely test whether it can be run affordably, operated day-to-day, or supported by existing staff and rules. Accordingly, pilots should test institutional readiness as much as technical feasibility.

4. Regulatory Pace Should be in Sync with AI Development

   Participants highlighted that governance systems, including data access rules, procurement processes, and accountability frameworks, often lag behind AI experimentation. When governance does not keep pace, attempts to scale can amplify risk rather than public value.

5. Structured Use cases May Insulate Against Any Policy Reversals

   Speakers argued that structured use case frameworks can help governments compare interventions, manage risk, allocate resources, and avoid ad-hoc or politically driven deployments, particularly under conditions of uncertainty and capacity constraints.

Safe AI Use cases at Scale

This closed-door session ran parallel to a closed-door session titled "AI for All: Catalyzing Global South Compute and Global Commons." The "Safe AI-Use Cases at Scale" session examined how trust, safety, and accountability shaped the real-world adoption of artificial intelligence, particularly in high-impact sectors such as agriculture, healthcare, education, and governance. Moving beyond abstract debates on AI risk, the session focused on everyday safety challenges that arise when AI systems are deployed at population scale. The speakers in this session emphasized that safe AI is a socio-technical challenge, rooted as much in institutions, norms, and user contexts as in technical design, and highlighted India's experience with DPI as a valuable reference point for enabling trustworthy AI adoption in the Global South.

The primary objectives of the closed-door session were:

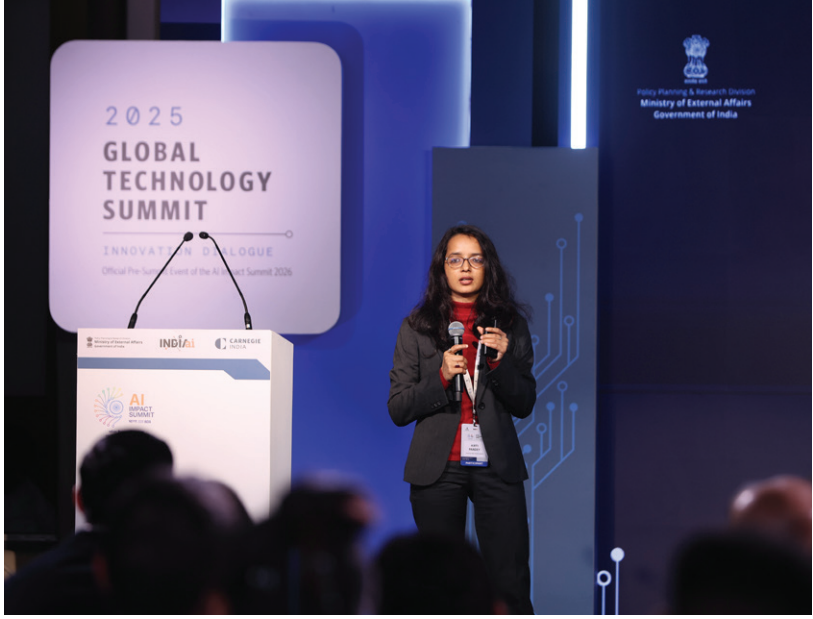

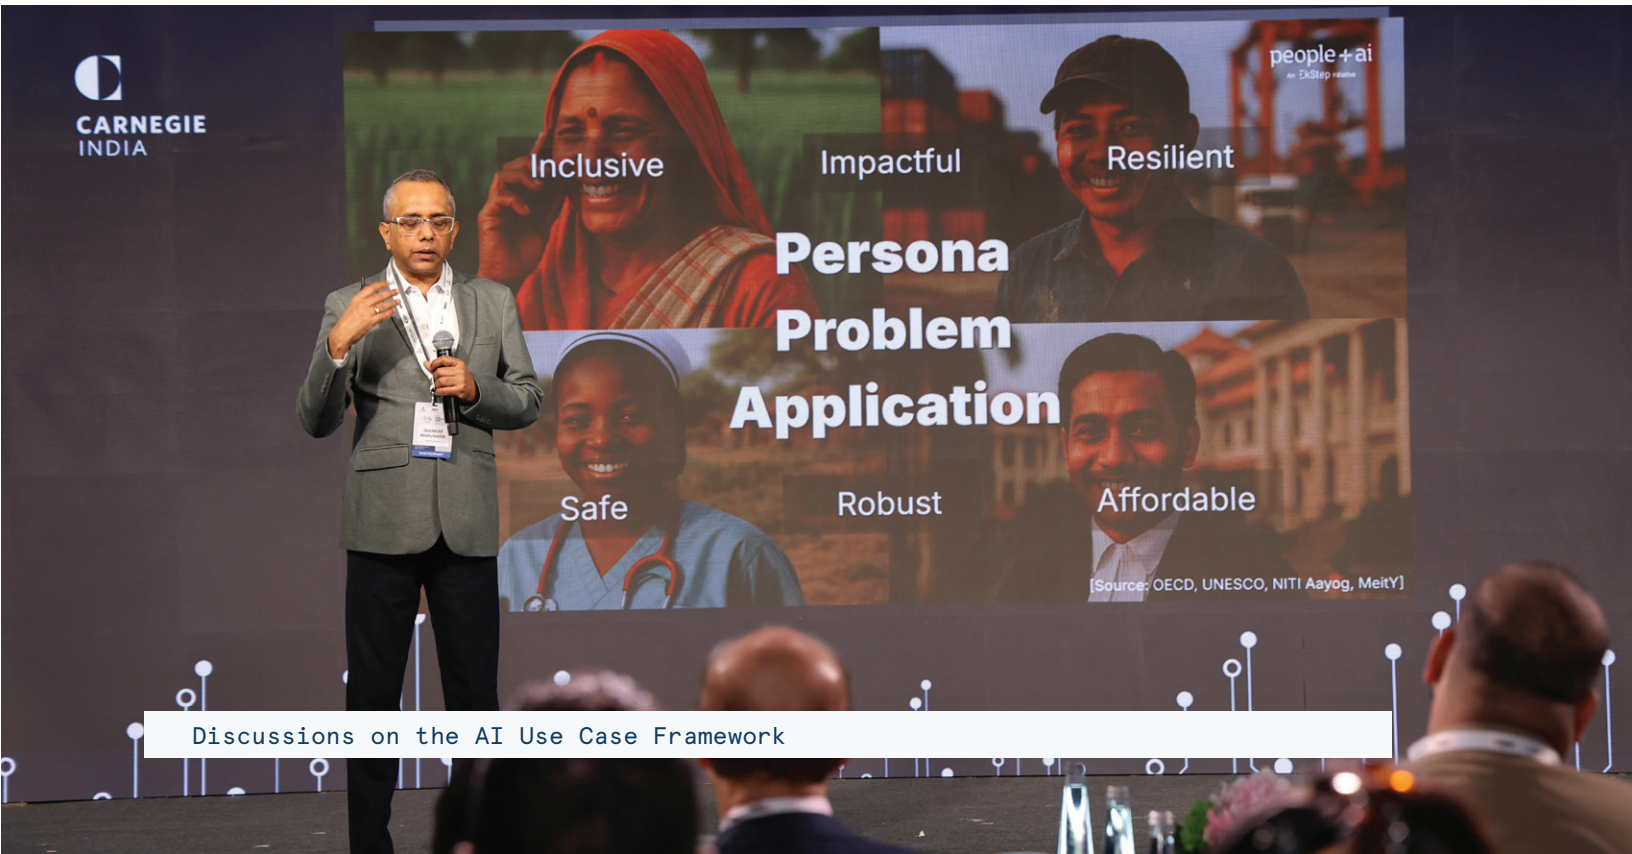

Discussions on the AI Use Case Framework

1. Trust and Safety as Foundational Requirements: To examine trust and safety as foundational requirements for large-scale AI adoption rather than post-deployment add-ons.
2. Ground Safety in Real-World Use Cases: To ground AI safety discussions in real-world use cases and sector-specific risks, particularly in high-stakes public domains.
3. Identify Governance and Bias Challenges at Scale: To identify governance, accountability, and bias-related challenges that emerge when AI systems operate at scale.
4. Explore Enabling Infrastructure for Safe AI: To explore the role of research, benchmarking, pilots, and digital public infrastructure in enabling safe and inclusive AI adoption.

Below are the key takeaways from this session:

1. Trust and Safety are Socio-Technical, Not Purely Technical Challenges

   Trust in AI systems, much like trust in UPI or other DPI, emerges through use, social acceptance, and institutional legitimacy rather than through model performance alone. Safety is experienced only when trust exists and is shaped by legal frameworks, social norms, enforcement mechanisms, and accountability structures. This makes trust and safety inherently contextual, layered, and dependent on institutions as much as on technology.

2. Safe AI Must Be Grounded in Concrete Use cases and User Contexts

   A strong consensus emerged against one-size-fits-all approaches to AI safety. Safety requirements vary significantly across sectors such as agriculture, healthcare, education, and governance, as well as across user groups with differing levels of literacy and vulnerability. While core principles, reliability, predictability, transparency, bias mitigation, and human accountability, remain constant, effective safety design must be people-first and anchored in specific use cases rather than abstract technical ideals.

3. Reliability, Accountability, and Bias Constitute the Primary Risks at Scale

   Inconsistent or incorrect AI behaviour was identified as the most immediate threat to trust and safety, particularly in high-stakes domains. Accountability gaps remain a critical weakness, as responsibility for AI outcomes is often diffused across developers, deployers, institutions, and regulators. Persistent representational biases, across gender, caste, religion, race, and region, were shown to be dynamic and difficult to eliminate, requiring explicit trade-offs and normative judgment rather than purely technical solutions.

4. Safety Must be Evaluated Across Full Systems, Not Isolated Models, or Prompts

   Evidence from real-world deployments, including large-scale agricultural AI systems, demonstrated that safety failures often emerge over extended interactions rather than single prompts. Risks such as conversational drift, contextual misinterpretation, and domain leakage highlight the need for end-to-end testing across entire user journeys. Both vertical (domain-specific) and horizontal (crosscutting) safety mechanisms must be stress-tested through iterative testing, sandboxes, and continuous user feedback from early stages of deployment.

5. Safe AI Adoption Depends on Real-World Experimentation, Local Collaboration, and Infrastructure

   Benchmarks and metrics remain essential for measuring progress, but they are inherently limited in capturing real-world complexity. Learning must therefore come from pilots, controlled experimentation, and iterative scaling. Local expertise, particularly in the Global South, is indispensable for designing relevant benchmarks, datasets, and evaluations. Open, interoperable digital public infrastructure can embed safety, accountability, and trust into AI systems at scale, enabling integration with existing public systems while avoiding duplication and concentration of power.

AI for All: Catalyzing a Global South Compute and Knowledge Commons

The global AI landscape is rapidly consolidating around a handful of nations and corporations with access to massive computational resources, creating an emerging "compute divide" that threatens to entrench new forms of technological dependency. Low-resource regions risk becoming perpetual consumers of AI technologies they cannot access, customize, or govern, undermining both their development trajectories and technological sovereignty. Without equitable access to compute infrastructure, countries lose the ability to train models that reflect their languages, cultural contexts, and policy priorities, their data flows outward and their capacity to shape AI-driven futures diminishes. This closed-door session convened senior leaders from government, industry, multilateral institutions, and philanthropy to confront this challenge.

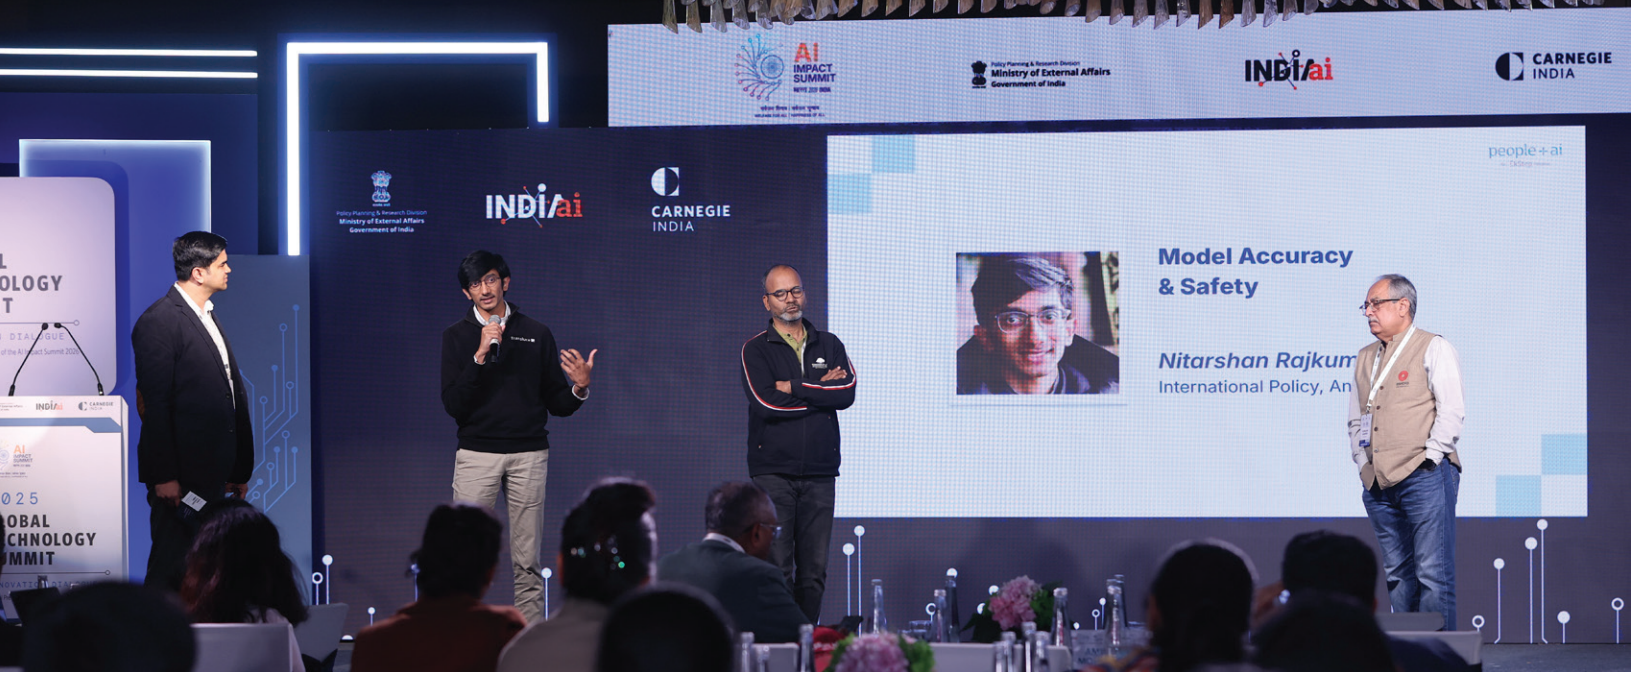

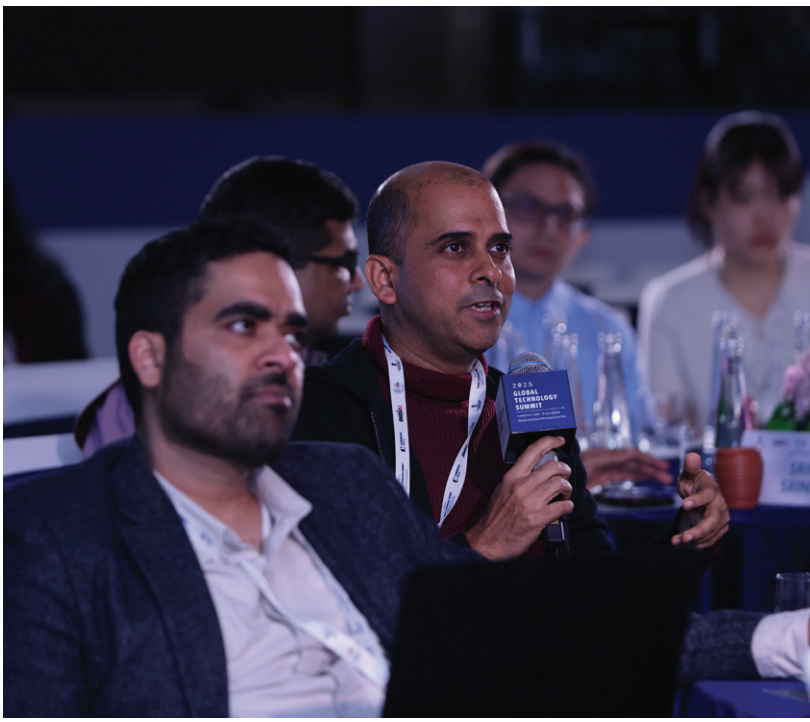

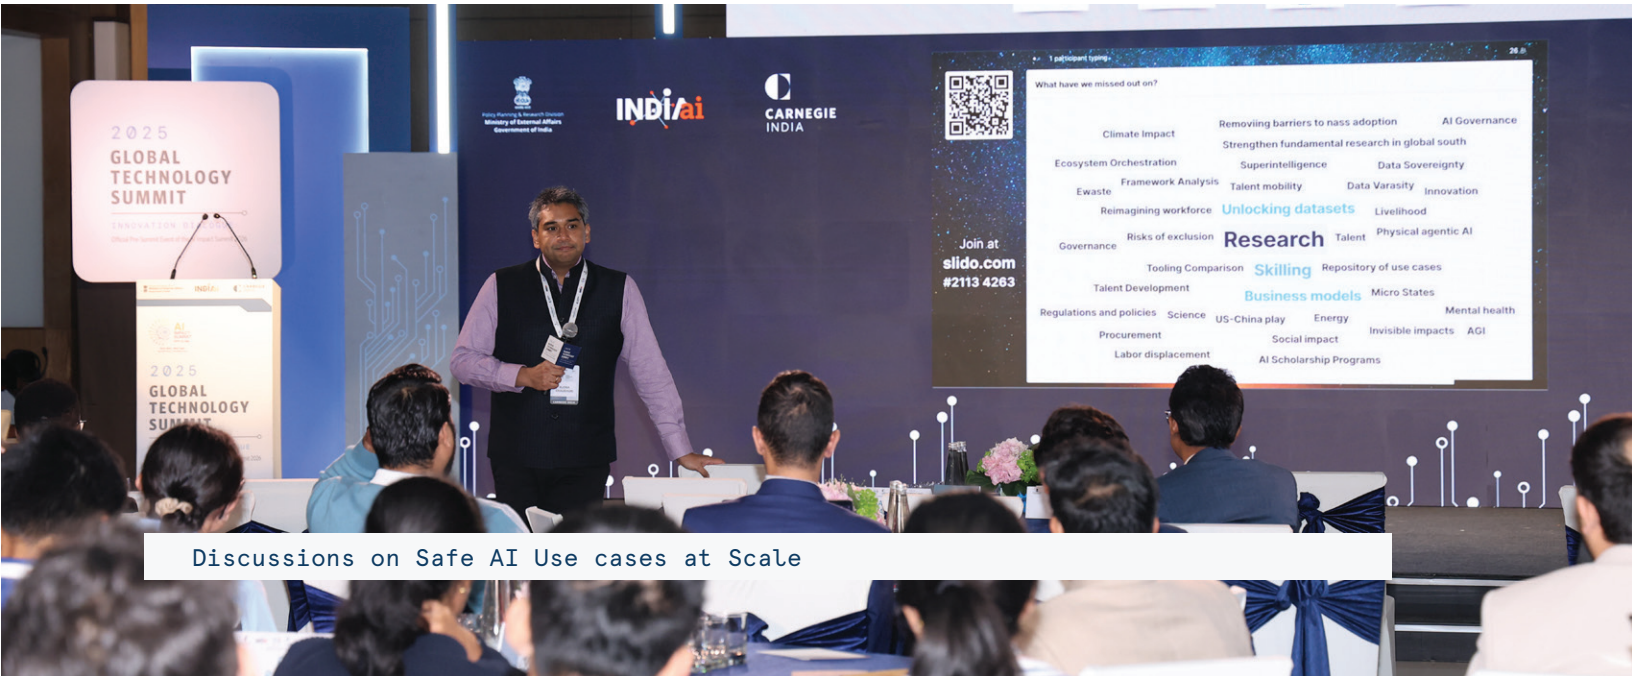

Discussions on Safe AI Use cases at Scale

The primary objectives of the closed-door session were:

1. Define the Urgency of Equitable Compute Access: Articulate why compute access is critical for achieving both development and sovereignty goals in low-resource regions, establishing the strategic imperative for addressing the compute divide before dependencies deepen and capacity gaps widen irreversibly.
2. Explore International Cooperation Models: Examine proven frameworks from global health and scientific research to identify replicable governance, coordination, and resource-sharing mechanisms applicable to a Global South AI Commons.
3. Outline Financing Pathways for Shared Infrastructure: Begin mapping concrete financing strategies capable of supporting shared compute and knowledge infrastructure at scale, including blended finance models, multilateral instruments, philanthropic partnerships, and innovative mechanisms that ensure sustainability and equitable access.

Below are the key takeaways from this session:

1. The Nature of the Compute Divide

   The global compute challenge is one of extreme concentration rather than absolute scarcity, with over half of the total capacity in the United States and less than 5 percent in the Global South. This imbalance means tasks that take months in the Global North could take decades elsewhere, reinforcing structural inequality. Treating compute as a single resource obscures important distinctions between training, inference, edge computing, and specialized workloads.

2. Use Case–Driven Infrastructure Needs

   Most high-impact applications do not require frontier-scale training infrastructure. Inference, fine-tuning, and edge deployments are far more relevant for many countries, especially in the Global South. Policymakers often prioritize expensive training infrastructure when most national needs are inference-based and far less compute-intensive.

3. Spectrum of AI Sovereignty

   AI sovereignty exists along a spectrum rather than as a binary choice. Countries face trade-offs between cost, complexity, and control, and will likely adopt hybrid strategies that vary by sector and sensitivity rather than pursuing full-stack self-sufficiency. AI infrastructure depends on power, data centers, applications, and chips, of which only chips are largely out of near-term reach for most countries. Focusing on the other three layers enables progress while avoiding oversimplified "GPU-only" strategies.

4. Energy and Sustainability as Binding Constraints

   Reliable power supply, cooling, and energy regulation are foundational to AI infrastructure. In many contexts, energy availability and carbon intensity are more constraining than access to chips themselves, raising sustainability and equity concerns.

5. Structural Dependencies and Ecosystem Lock-In

   Concentration in hardware and software ecosystems creates long-term dependencies. Open-source alternatives offer potential pathways to reduce lock-in but require sustained ecosystem-building and coordination.

6. Compute Is Necessary but Not Sufficient

   Compute infrastructure alone does not guarantee adoption or impact. Utilization depends on complementary enablers such as data availability, skills, institutional capacity, developer tools, and affordable access mechanisms.

7. Financing and Sequencing Challenge

   Different stages of AI ecosystem development require different financing approaches. Grants and concessional finance are critical early on, while commercial investment becomes viable only once demand, use cases, and utilization pathways are established.

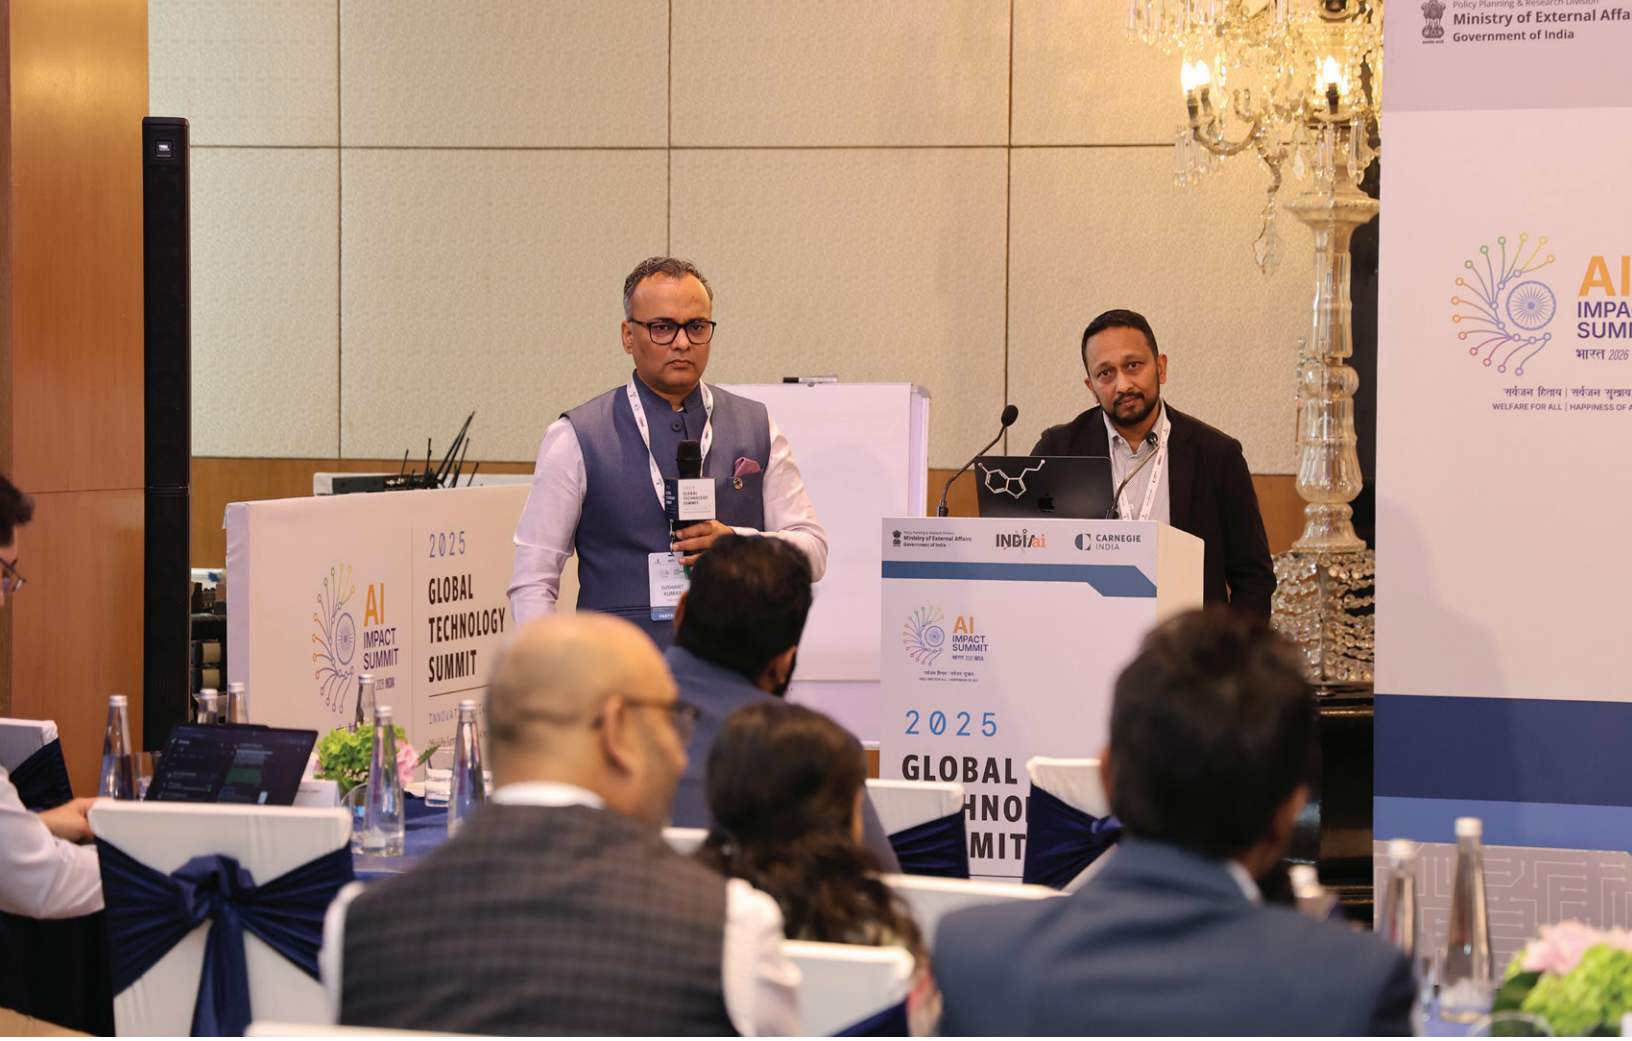

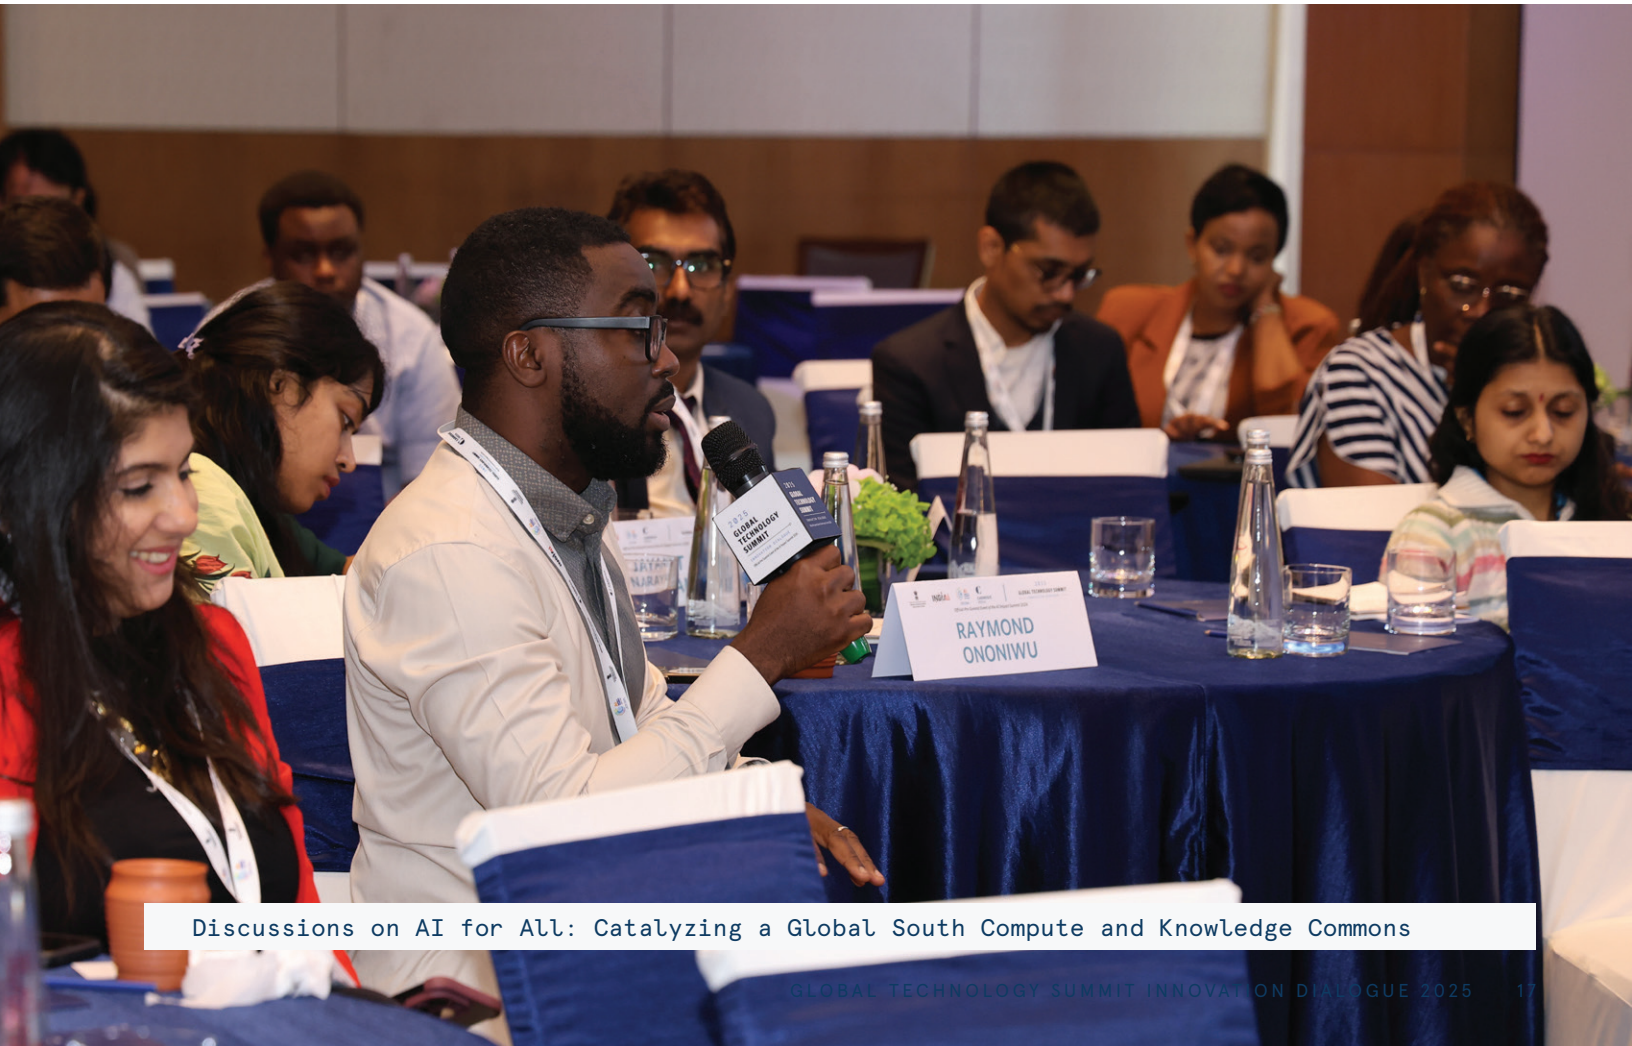

Discussions on AI for All: Catalyzing a Global South Compute and Knowledge Commons

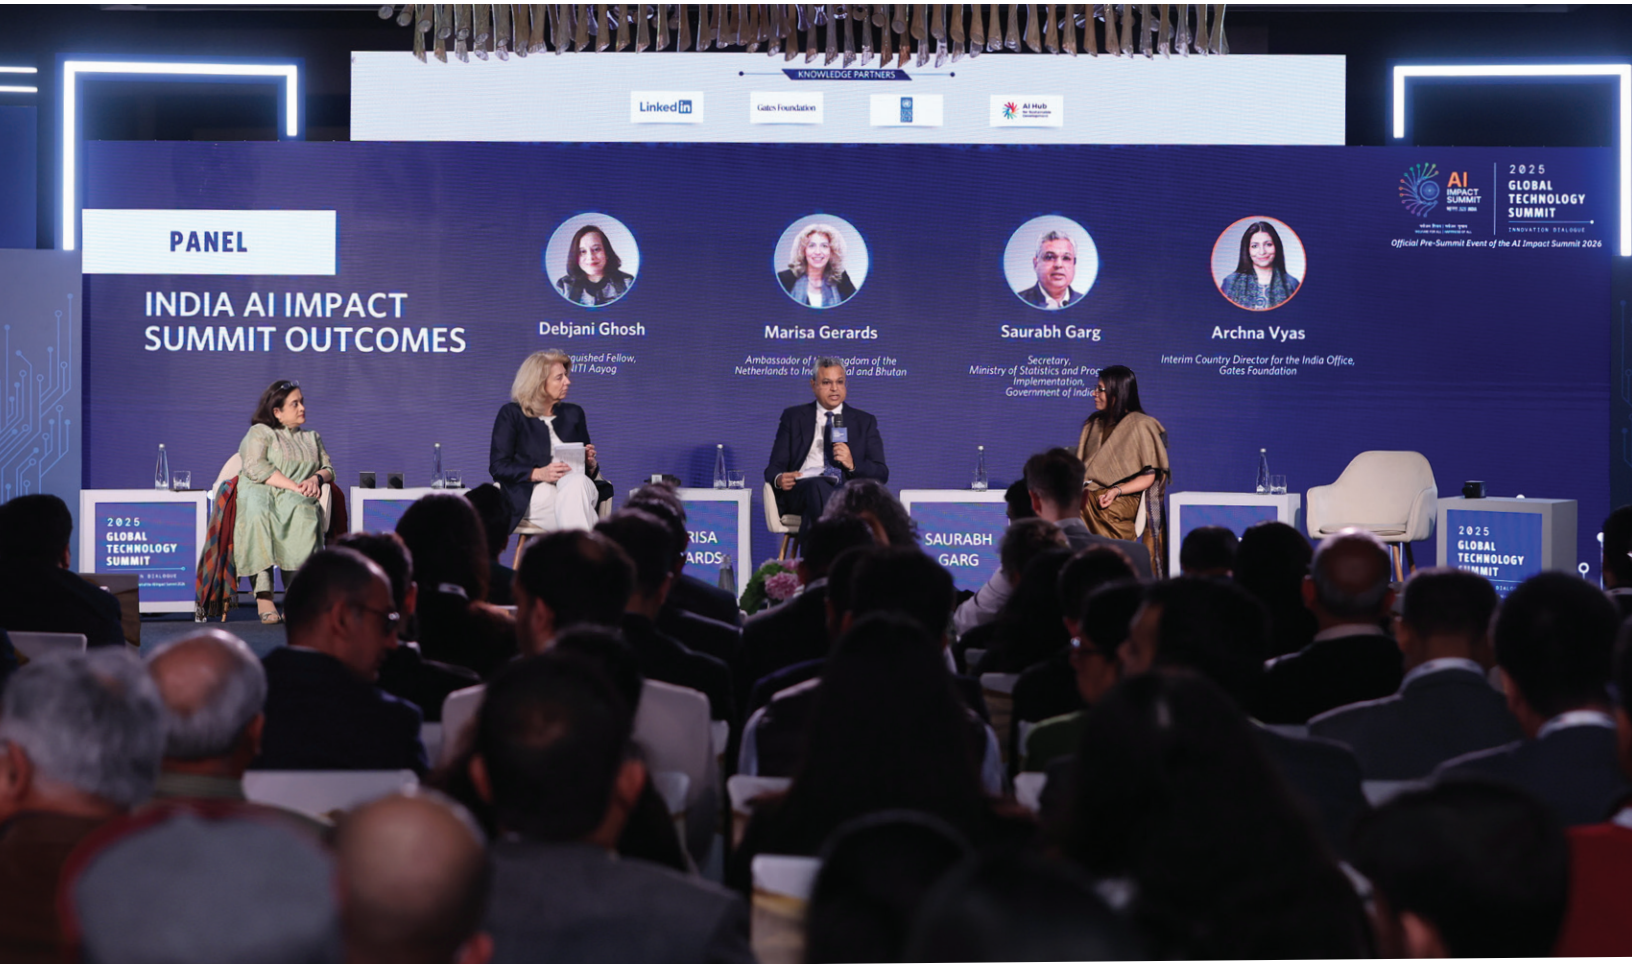
KNOWLEDGE PARTNERS
PANEL
INDIA AI IMPACT SUMMIT OUTCOMES
Debjani Ghosh
Marisa Gerards
Saurabh Garg
Archna Vyas
Interim Country Director for the India Office, Gates Foundation
2025 GLOBAL TECHNOLOGY SUMMIT
Official Pre-Summit Event of the AI Impact Summit 2026

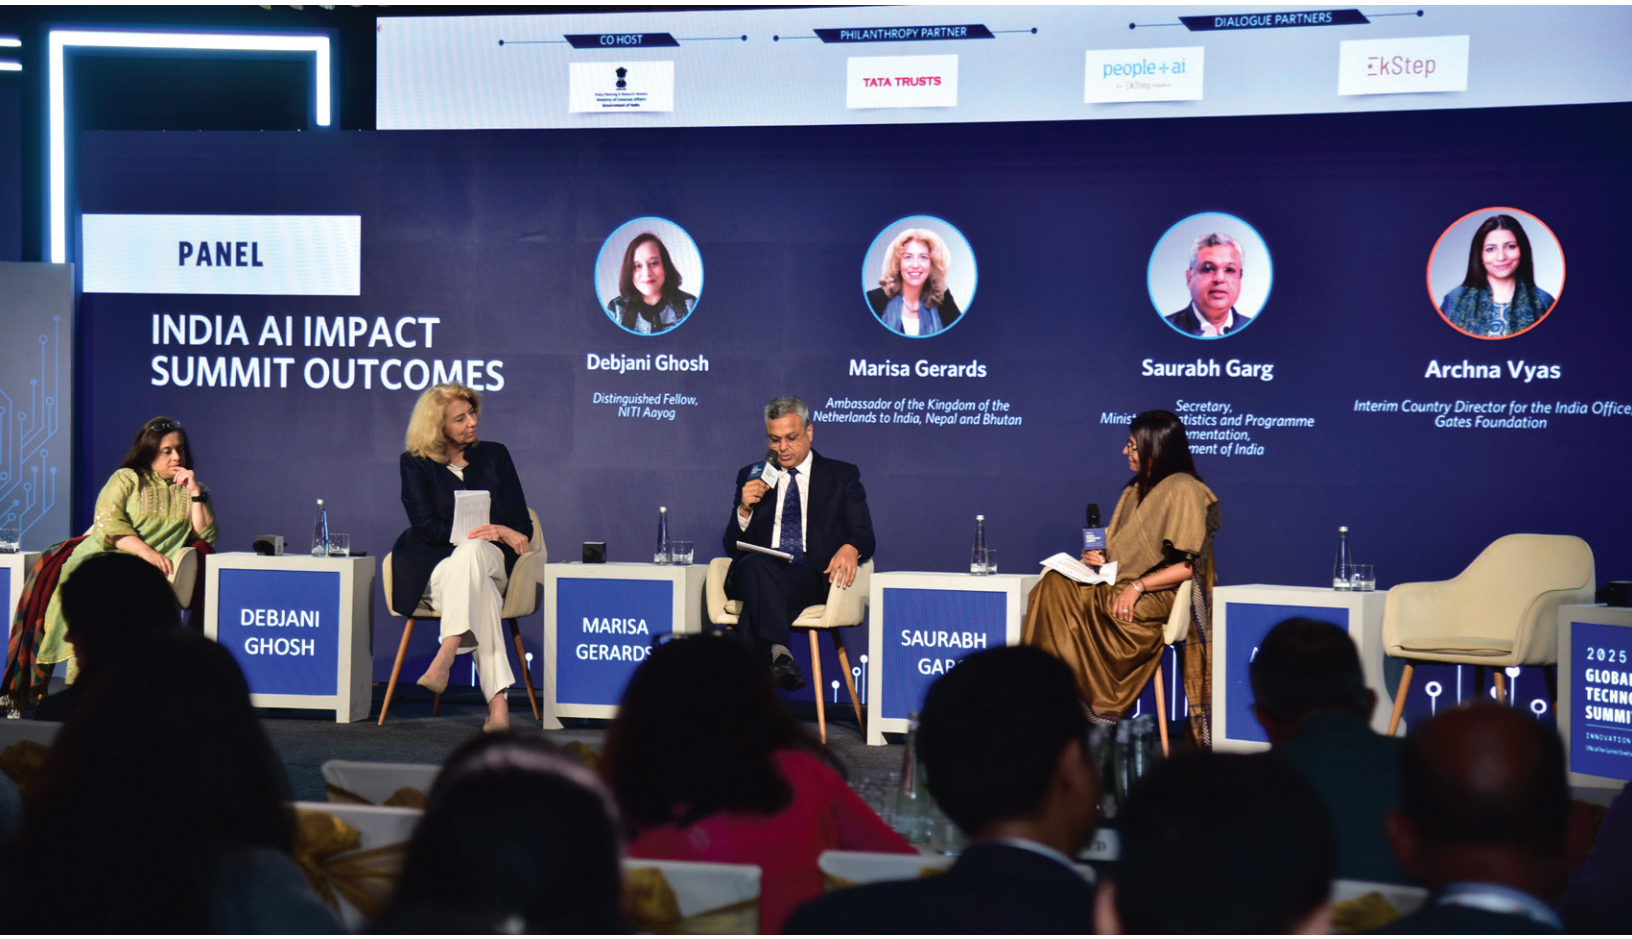
CO HOST
PHILANTHROPY PARTNER
TATA TRUSTS
DIALOGUE PARTNERS
people+ai
EkStep
PANEL
INDIA AI IMPACT SUMMIT OUTCOMES
Debjani Ghosh
Distinguished Fellow, NITI Aayog
Marisa Gerards
Ambassador of the Kingdom of the Netherlands to India, Nepal and Bhutan
Saurabh Garg
Archna Vyas
Interim Country Director for the India Office
Gates Foundation
DEBJANI GHOSH
2025 GLOBAL TECHNOLOGY SUMMIT

PUBLIC PANELS

India AI Impact Summit Outcomes

The opening public panel of the GTS Innovation Dialogue brought together leaders from India's international partners to discuss the vision, priorities, and expected outcomes of the India AI Impact Summit 2026. Co-chairs of two critical working groups, "Democratizing AI Resources" and "AI for Economic Growth and Social Good," were joined by The Netherlands' Ambassador to India to discuss the thematic focus of the AI Impact Summit 2026, the reason India pivoted from "action" towards thinking about "impact," and the opportunities and challenges that the AI Impact Summit 2026 has presented.

The primary objectives of the public panel were:

1. Introduce the AI Impact Summit Framework: Present the vision, structure, and seven working groups driving the India AI Impact Summit 2026, emphasizing the deliberate shift from discussing AI's potential to measuring its actual impact on economic and social development.
2. Articulate Working Group Priorities: Detail the specific focus areas of key Working Groups, including democratizing access to AI resources (compute, datasets, models, and talent) and accelerating pathways to impact creation across diverse countries regardless of their development stage.
3. Foster International Collaboration: Highlight the role of global partnerships, blended financing mechanisms, and multi-stakeholder engagement (governments, international organizations, private sector, academia, philanthropies) in scaling AI solutions from pilots to impact.
4. Establish Success Metrics and Legacy: Outline concrete deliverables, including the impact repository, global recognition frameworks that recognize and reward successful AI stories, best practice sharing mechanisms, and sustainable collaborative consortium that will outlast the summit itself.

Below are the key takeaways from this public session:

1. Reframing the Global AI Agenda Around Impact

   The panel underscored that India took a deliberate step to shift the discourse from the abstract promises and potential of AI toward demonstrable, measurable socioeconomic impact, particularly in areas such as livelihoods, healthcare, education, and agriculture. The panel highlighted that especially for countries in the Global South, it is context-specific solutions with the potential to scale that will bring impact rather than frontier AI research.

2. Democratizing Access to AI Resources Is Foundational for Equity

   Equitable access to core AI resources, including compute, data, models, and talent, is a key prerequisite for inclusive AI development. The panel strongly emphasized the need for affordability, agency, and reducing structural concentration of these resources. The focus of the working group on democratizing AI resources is therefore, developing concrete mechanisms to bolster access, including a compute discovery and access platform, multilingual and multicultural datasets, and domain-specific models for edge deployment.

3. Scaling Impact Requires Moving Beyond Pilots to System-Level Deployment

   A recurring theme was the need to design AI solutions for scale from the outset, drawing lessons from DPI and emphasizing deployment pathways that translate innovation into population-level outcomes.

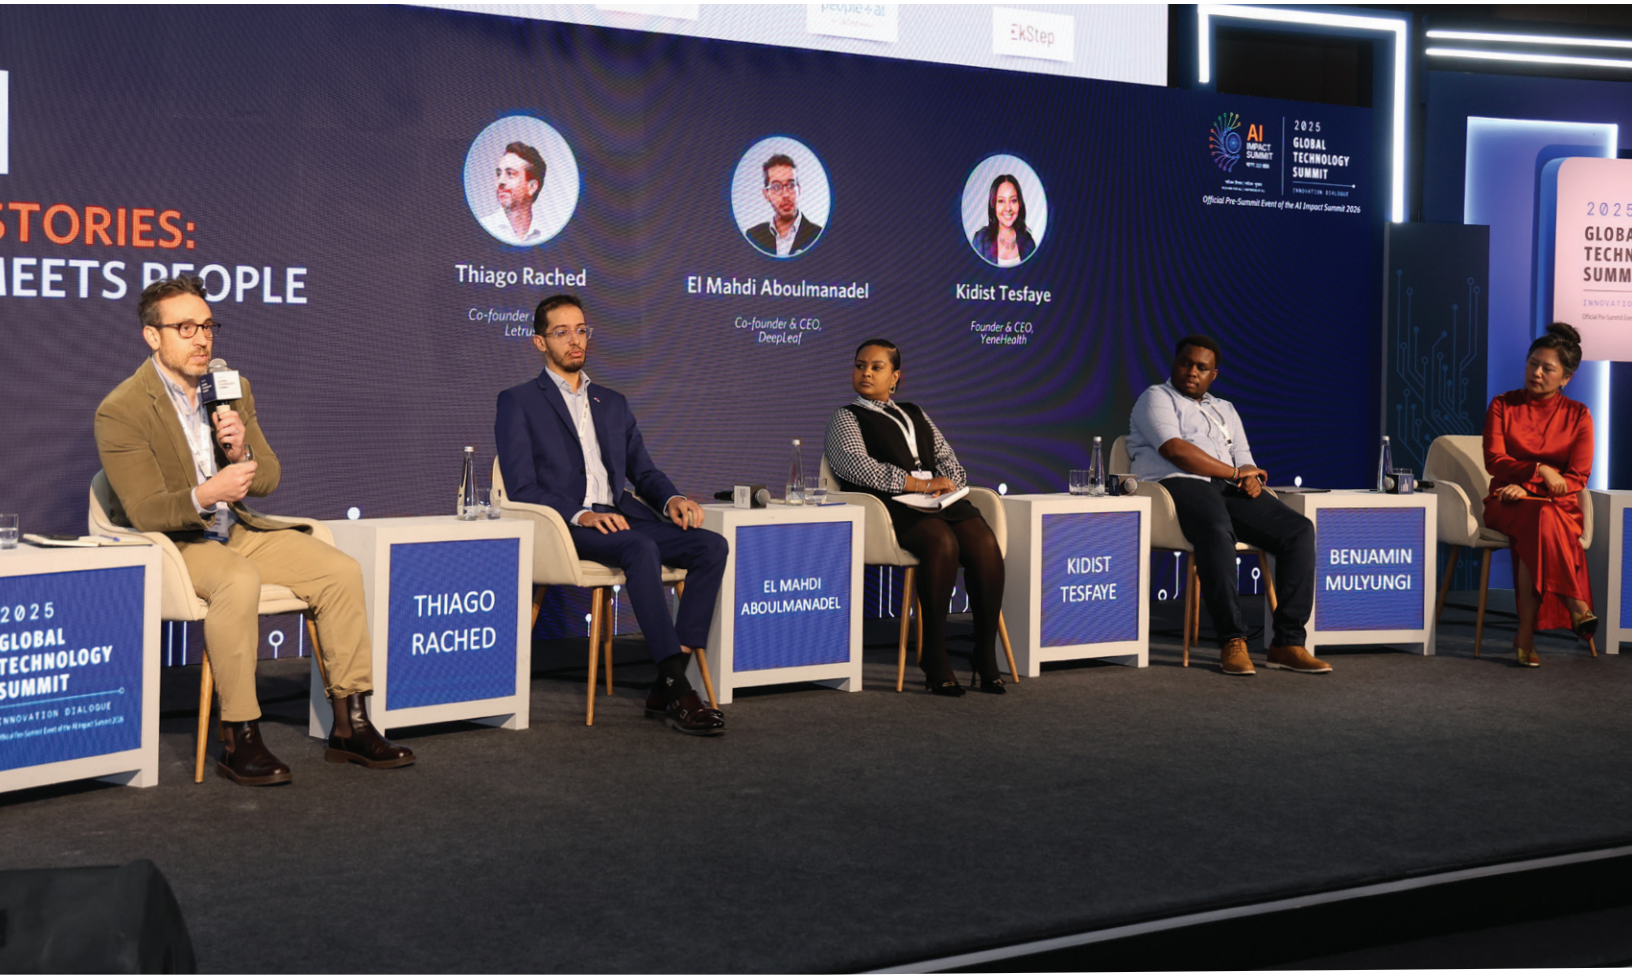
Thiago Rached
Co-founder
El Mahdi Aboulmanadel
Co-founder & CEO,
DeepLeaf
Kidist Tesfaye
Founder & CEO,
YeneHealth
THIAGO RACHED
EL MAHDI ABOULMANADEL
KIDIST TESFAYE
BENJAMIN MULYUNGI
2025 GLOBAL TECHNOLOGY SUMMIT
Official Pre-Summit Event of the AI Impact Summit 2026


AI IMPACT SUMMIT
Ministry of External Affairs
Government of India
INDIAai
CARNEGIE INDIA
PANEL
AI IMPACT STORIES:
WHERE AI MEETS PEOPLE
Benjamin Mwalimu
Head of Technology,
Keyzom Ngodup Massally
Director,
UNDP

4. Collaboration Must Be Institutionalized to Sustain Impact

   Beyond the actual summit, success depends on creating lasting processes for knowledge sharing, establishing impact measurement frameworks that mean the same thing globally, and building ecosystems where countries continuously contribute solutions, share learnings, and co-develop responses to development challenges. The panel discussion highlighted the importance of including all relevant stakeholders including governments, industry, academia, philanthropy, and international organizations, to enable knowledge-sharing, coordinated investment, and replication of proven impact models across countries.

AI Impact Stories: Where AI Meets People

This public panel examined how practitioners generate positive socio-economic impact through AI-enabled interventions in education, agriculture, and women's health across the "Global South." Speakers showed that AI use cases emerge not from technology deployment alone, but from clearly defined problem statements rooted in real service delivery contexts. They argued that impact develops when practitioners embed AI as a support layer within services, workflows, and frontline decision-making, rather than deploying it as a standalone system. The discussion also highlighted how institutional and governance constraints continue to shape whether these interventions move beyond pilots toward broader adoption, emphasizing the central role of collaboration and capacity in sustaining impact.

The primary objectives of the public panel were:

1. Understand How AI Use Cases Emerge: To examine how AI use cases take shape in practice across sectors and what factors influence their initial design and direction. Drawing on experiences from education, agriculture, and women's health, the session explored how practitioners identify priority problems and translate them into AI-enabled interventions that operate within existing service environments.
2. Identify Barriers to Moving Beyond Pilots: To examine why many AI use cases struggle to progress beyond pilot stages, even when they demonstrate early value. The discussion focused on recurring constraints that limit adoption, including access to usable data, the affordability and governance of compute, linguistic and cultural adaptation, and the capacity of institutions to absorb and sustain new technologies.
3. Assess Conditions for Sustained Impact: To assess what conditions enable AI-enabled interventions to deliver sustained socio-economic outcomes over time. The discussion examined how trust, human oversight, service design, public-sector buy-in, and regulatory clarity shape whether ministries, public service providers, and implementing partners can adopt and maintain AI use cases.

Below are the key takeaways from this public panel:

1. Use Cases Emerged as a Response to Gaps in State Capacity

   Across education, agriculture, and women's health, practitioners built AI use cases in response to concrete state service delivery gaps rather than technical ambition. In each sector, they applied AI to support human capacity, whether by augmenting feedback processes, improving decision support, or expanding access to information and care.

2. Pilots Fail in Absence of Government Support

   Efforts to move beyond pilots stalled less because models underperformed and more because AI tools had to operate within government systems not designed for data-driven automation. The speakers during this public session pointed to procurement processes, reliance on government-owned data, limited frontline capacity, and rigid service workflows as key barriers.

3. Compute Costs Can Be Managed with a Clearer Roadmap

   Compute constraints shaped both the feasibility and sustainability of AI deployment. While cloud credits enabled early experimentation, long-term use depended on alignment with government budgets, national infrastructure choices, and clear decisions about who bears recurring costs.

4. AI Needs to Augment, and Not Replace Humans to Build Trust

   Trust played a decisive role in adoption. In health and education contexts, panelists emphasized the need for human-in-the-loop oversight, careful interaction design, and strong data safeguards to ensure users viewed AI-supported decisions as legitimate and appropriate.

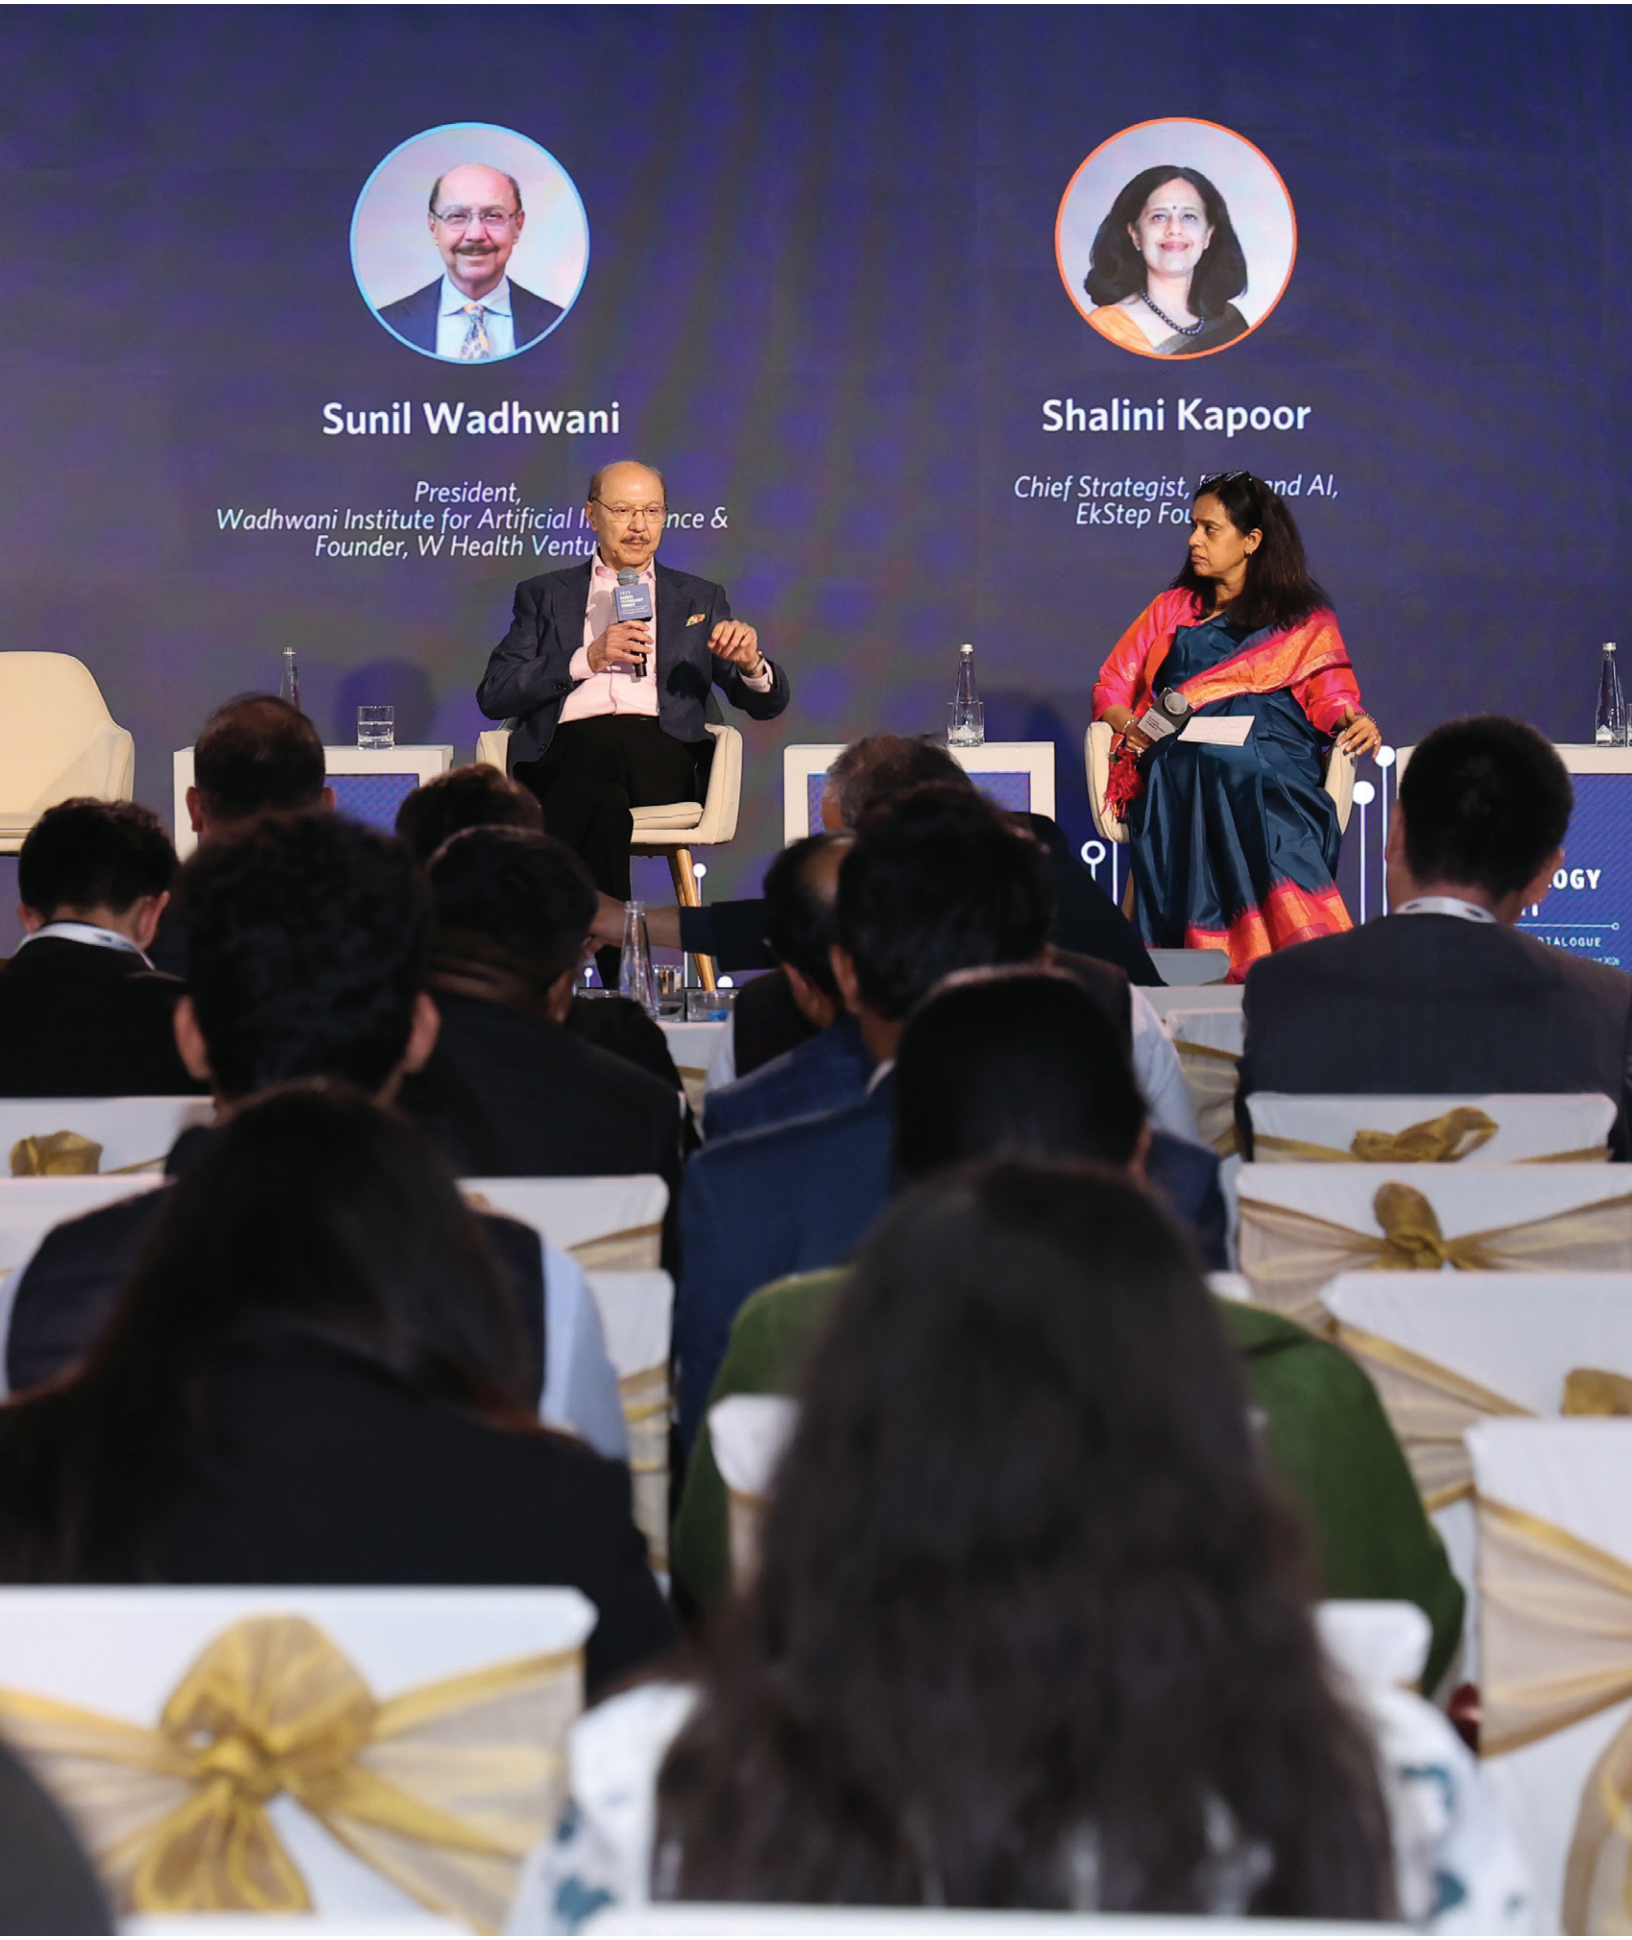
Sunil Wadhwani
President,
Founder, W Health Ventu
Shalini Kapoor
Chief Strategist,
EkStep Fou

5. AI Growth is Seldom Linear and Must Address the Needs of the Whole Ecosystem to Scale

   The discussion convened challenged assumptions that AI adoption follows a straight path from pilots to scale. Speakers described scale as a gradual process of institutionalization that requires regulatory approval, service integration, and incentive alignment across ministries, providers, and users.

Conversation with Sunil Wadhwani: AI for the Last Mile

This conversation with Sunil Wadhwani offered a compelling vision for how AI could drive inclusivity across the Global South.

The primary objectives of the Conversation were:

1. To Bridge the Digital Divide: The conversation wanted to explore how AI can be deployed in Tier-4 cities and the Global South with the same efficacy as in Silicon Valley in the U.S.
2. To Identify Scaling Bottlenecks: To analyze why most AI applications fail to reach mass adoption and how to overcome implementation hurdles.
3. To Look at India's Position as a Global Leader: To gauge whether India's AI and digital infrastructure can serve as a blueprint for other developing nations in South Asia, Africa, and Latin America.
4. To Promote System-Level Thinking: To move beyond individual apps to create national AI roadmaps involving data governance, capacity building, and policy.

Below are the key takeaways from this Conversation:

1. Addressing the AI Implementation Gap

   The speaker asserted that the technical model of AI is only 30 percent of the solution. The remaining 70 percent is the "last-mile" challenge of diffusing the application to enable impact. To solve this implementation gap and last-mile challenge, it was emphasized that a "workflow integration" was necessary, noting that AI fails when it's considered just as an "add-on." It was also noted that AI must be woven into the existing daily habits of frontline workers. Some frontline realities must also be contended with. Solutions must work in low-connectivity areas, support local languages, and be simple enough for non-technical users like primary school teachers.

   Examples were given of Wadhwani's AI application in education and health, and how it assists in reducing the drop-out rates in early education by promoting proficiency in their mother tongue.

2. India's lead on the "Global South Blueprint" on AI.

   It was noted that India is positioned as a leader in AI because it has already solved for population-scale digital identity (Aadhaar) and payments (UPI). Furthermore, AI was seen as the next layer of the DPI. India's AI approach was state-supported but privately innovated, focusing on public goods like health and education, as opposed to being purely commercially driven. In the U.S, AI development is led by larger technological companies that don't think as much about the last-mile impact of their models or solutions. Additionally, China, through its state-led model, has a different perspective on AI use cases. By building AI that plugs into a national health database, solutions can scale to hundreds of millions of people instantly rather than staying stuck in localized pilots. The India AI Impact Summit 2026 provides a promising opportunity to lead the AI footprint for all Global South countries that face constraints on resources and can adopt a mix of public-private models to address various challenges.

3. Concerns of Job Displacement

   The conversation had a stark warning of AI displacing jobs. Through AI, entry-level roles in law, accounting, and coding are being automated faster than new jobs are being created. It was observed that there is near unanimity on "reskilling" as there is a lack of knowledge on what the jobs of the future will look like. A shift toward AI supervision, data ethics, and analytics as core survival skills was urged.

4. Sovereignty Through Ecosystems

   For AI to be inclusive, it must be locally grown. "Importing" models from the West leads to problems like data bias, where Western AI models don't understand the nuances of a Tier-4 city or a village in South Asia, Latin America, or Africa. Without a local academic and research ecosystem, the Global South will remain "consumers" rather than "creators." Massive investment was called for in South Asian and African AI research centers to ensure the technology reflects local values and data.

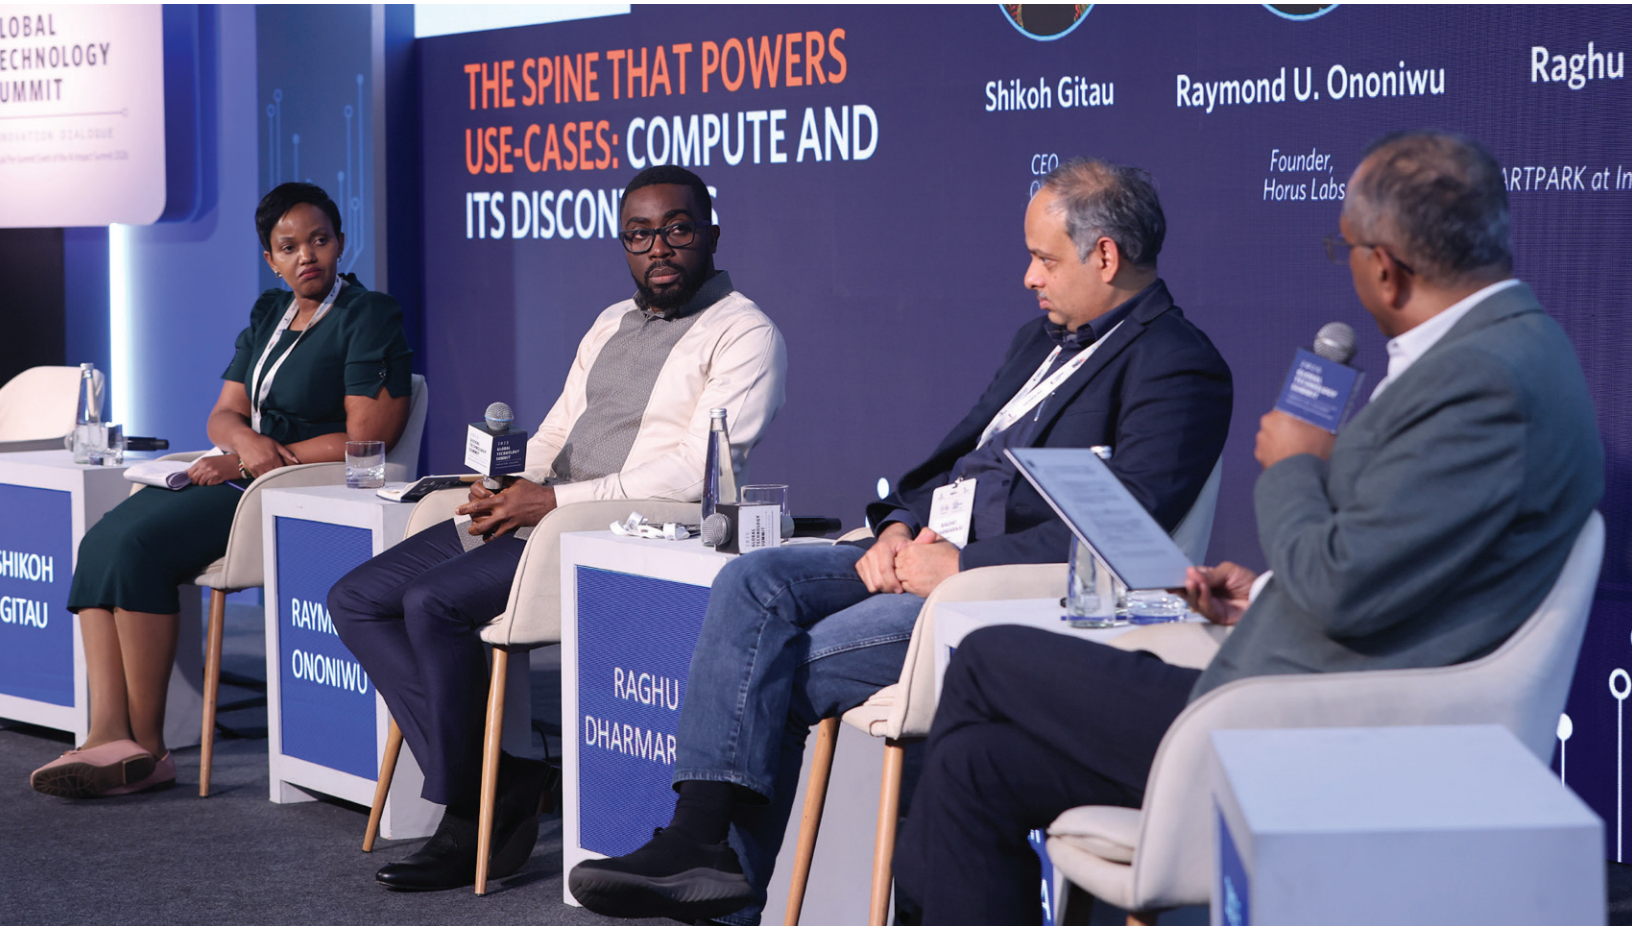
THE SPINE THAT POWERS
USE-CASES: COMPUTE AND
Shikoh Gitau
Raymond U. Ononiwu
Raghu
Founder,
Horus Labs
RAGHU
DHARMAR

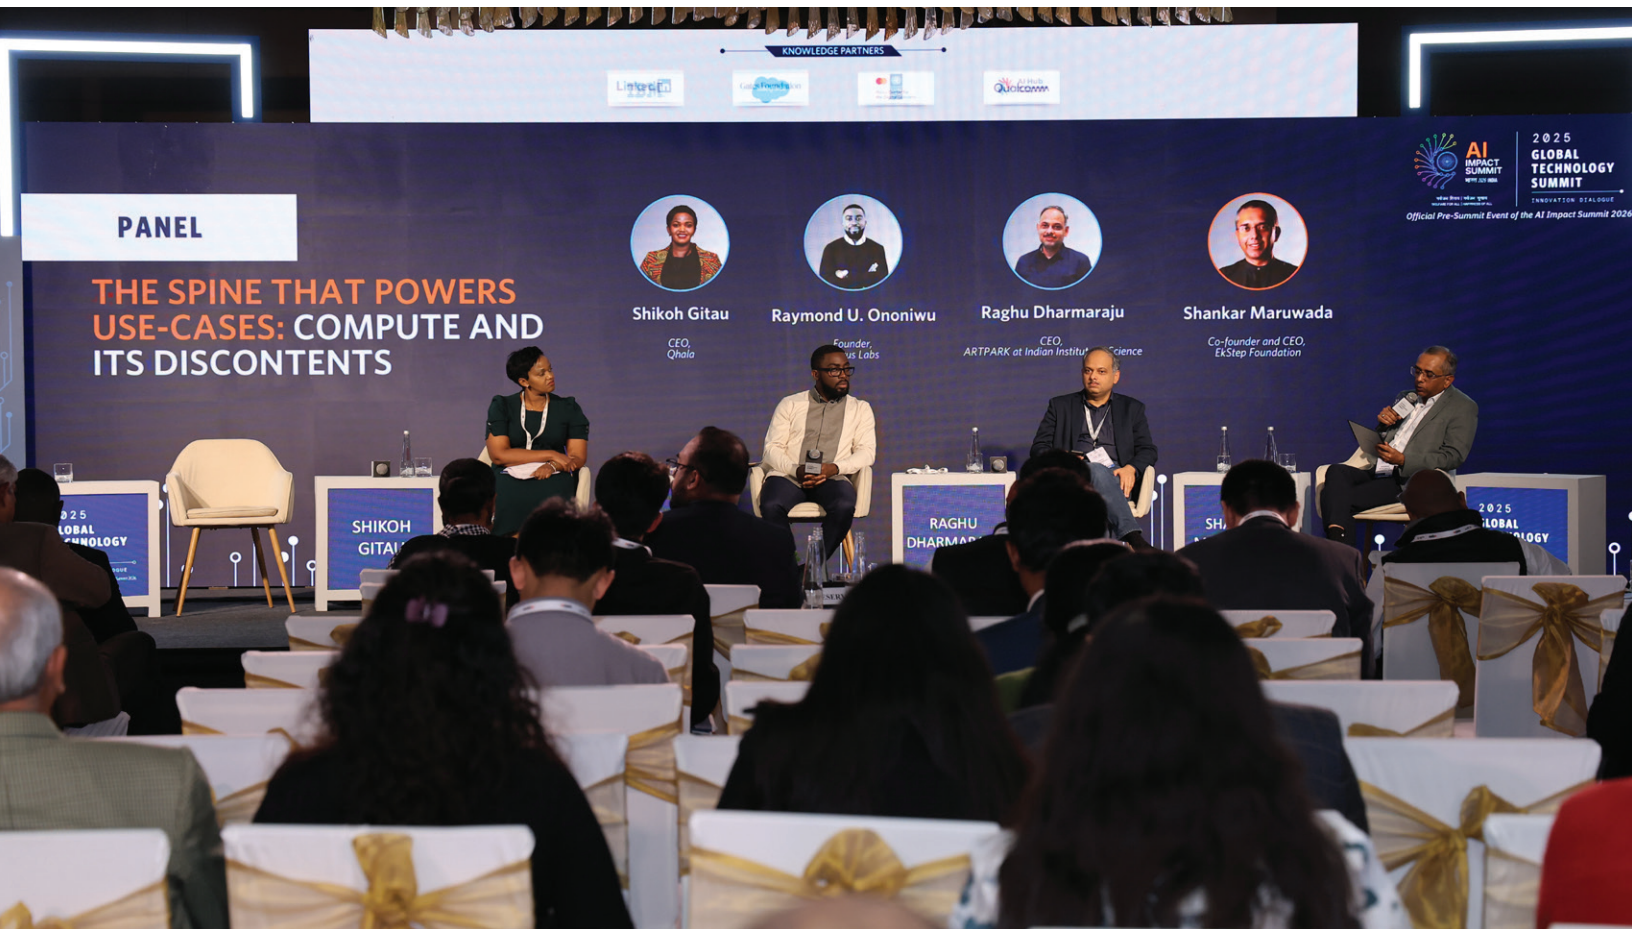
KNOWLEDGE PARTNERS
PANEL
THE SPINE THAT POWERS
USE-CASES: COMPUTE AND
ITS DISCONTENTS
Shikoh Gitau
CEO,
Qhala
Raymond U. Ononiwu
Raghu Dharmaraju
Shankar Maruwada
Co-founder and CEO,
EkStep Foundation
AI IMPACT SUMMIT
2025
GLOBAL
TECHNOLOGY
SUMMIT
INNOVATION DIALOGUE
Official Pre-Summit Event of the AI Impact Summit 2026
SHIKOH
RAGHU

The Spine that Powers Use cases: Compute and its Discontents

The public panel titled "The Spine that Powers Use cases: Compute and its Discontents" brought together practitioners from Africa and India to examine the realities of scaling AI solutions, including managing compute costs, accessibility, and infrastructure gaps in diverse environments.

The primary objectives of the public panel were as follows:

1. Interrogate Compute as a Critical Enabler: To interrogate compute as a critical enabler of AI solutions, moving beyond pilots to population-scale deployment.
2. Balance Sovereignty, Affordability, and Resilience: To clarify how governments and innovators can balance sovereign, edge, and cloud-based models while addressing affordability and resilience.

Below are the key takeaways from this public panel discussion:

1. Compute Sovereignty Should Be A Priority For All

   Compute must be a national priority for all countries alike. When nations fail to build their own sovereign infrastructure, they risk becoming permanent consumers of the models, laws, and priorities of other locales. Sovereignty is essential not only for national security but also for long-term technological resilience and independence.

2. Trade-Offs Will Have to Be Made by Global South Countries

   The Global South will need to adopt a different paradigm for AI development; one that is values-based and utility-oriented rather than capital-driven or control-focused. Regions such as India and Africa, with their diversity and scale, are uniquely positioned to lead in value-based AI development that prioritizes social utility and inclusion.

3. The Discussion on the Magnitude of Compute Required is a Nuanced One

   Questions surrounding long-term compute infrastructure cannot be reduced to binaries of centralized versus decentralized or sovereign versus cloud dependent. Compute is not a monolith but a stack of several layers that extends beyond GPUs to include energy, talent, datasets, and markets. It must be understood as an enabler of general-purpose technology. Countries in the Global South were encouraged to ask themselves the fundamental question of "why compute?" when designing strategies, ensuring that infrastructure decisions are tied to real use cases and societal needs.

4. Limited Access to Essential Elements of AI Development Can Be Managed

   Barriers such as data quality, the fragmented nature of regulations, and limited access to compute were identified as constraints to scale. The discussion highlighted the importance of meeting innovators where they are, including through sachet-style access models that mirror mobile payment systems, enabling hobbyists and start-ups to access compute affordably and flexibly.

5. Compute Demand Is Variable

   The discussion underscored the need to design compute strategies around specific use cases. Not all applications require maximum accuracy or large-scale models; trade-offs between cost, accuracy, and impact must be carefully considered. Techniques such as cache-ing, compression, and edge deployment can significantly reduce compute demands. AI was described as only one part of broader solutions, with much of the value lying in non-AI processes and human systems.

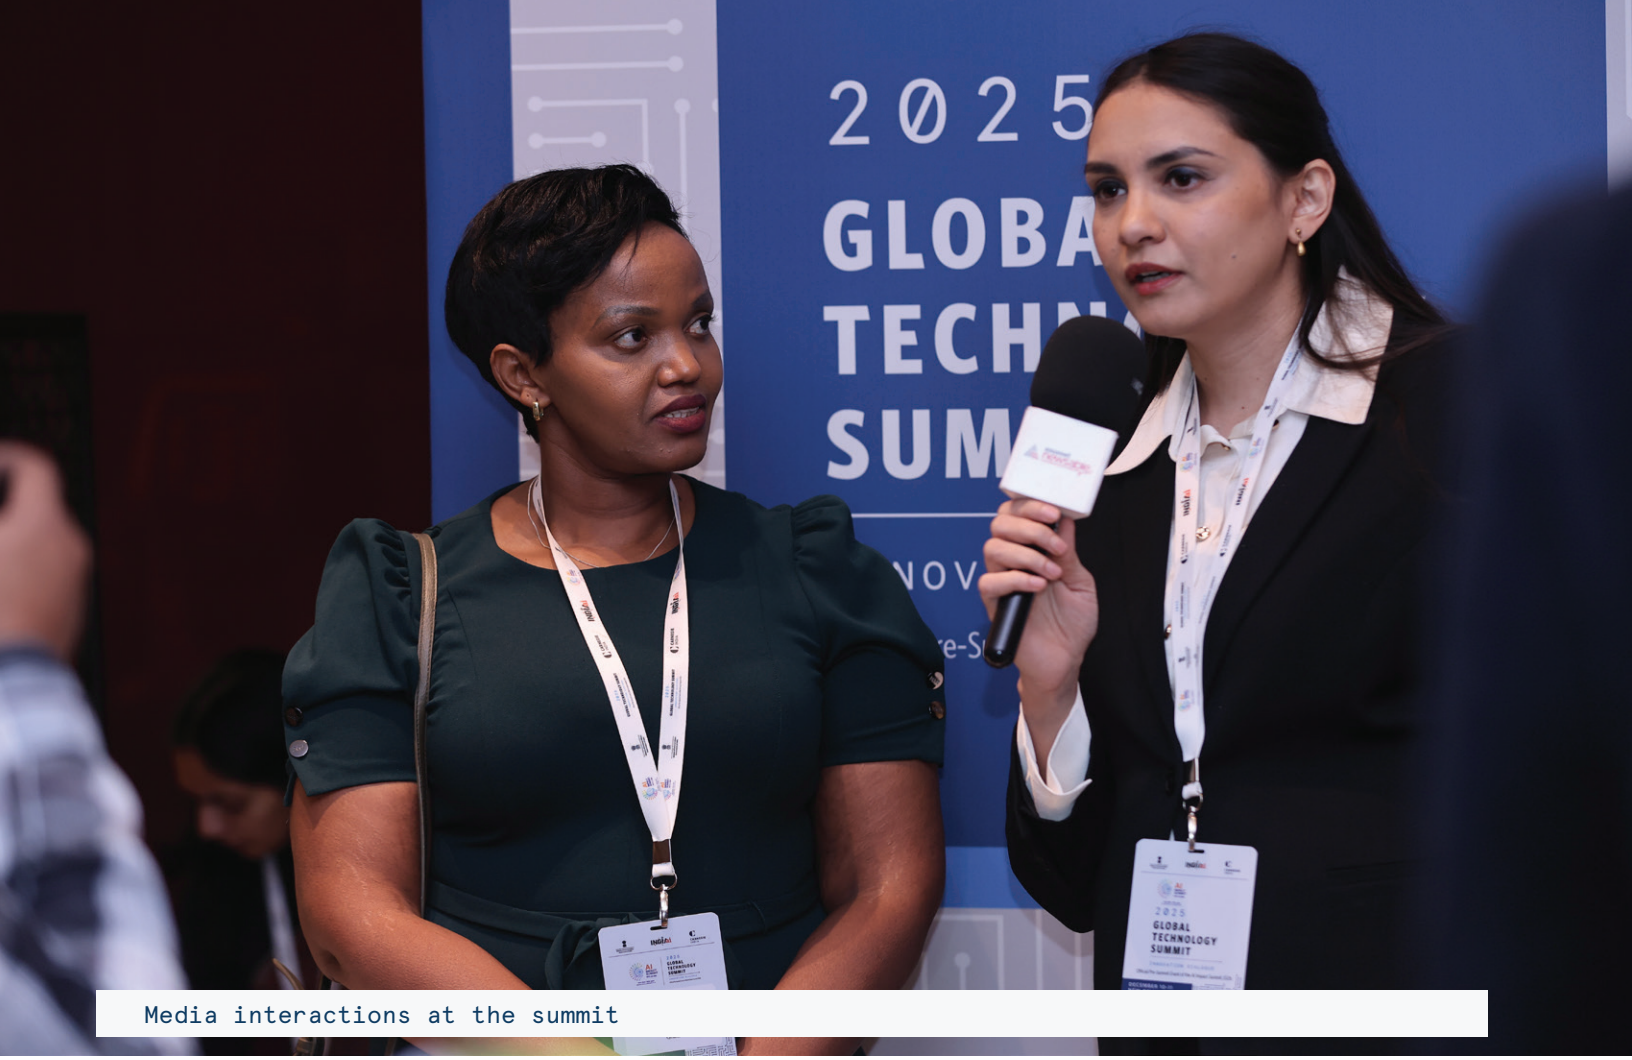

Media interactions at the summit

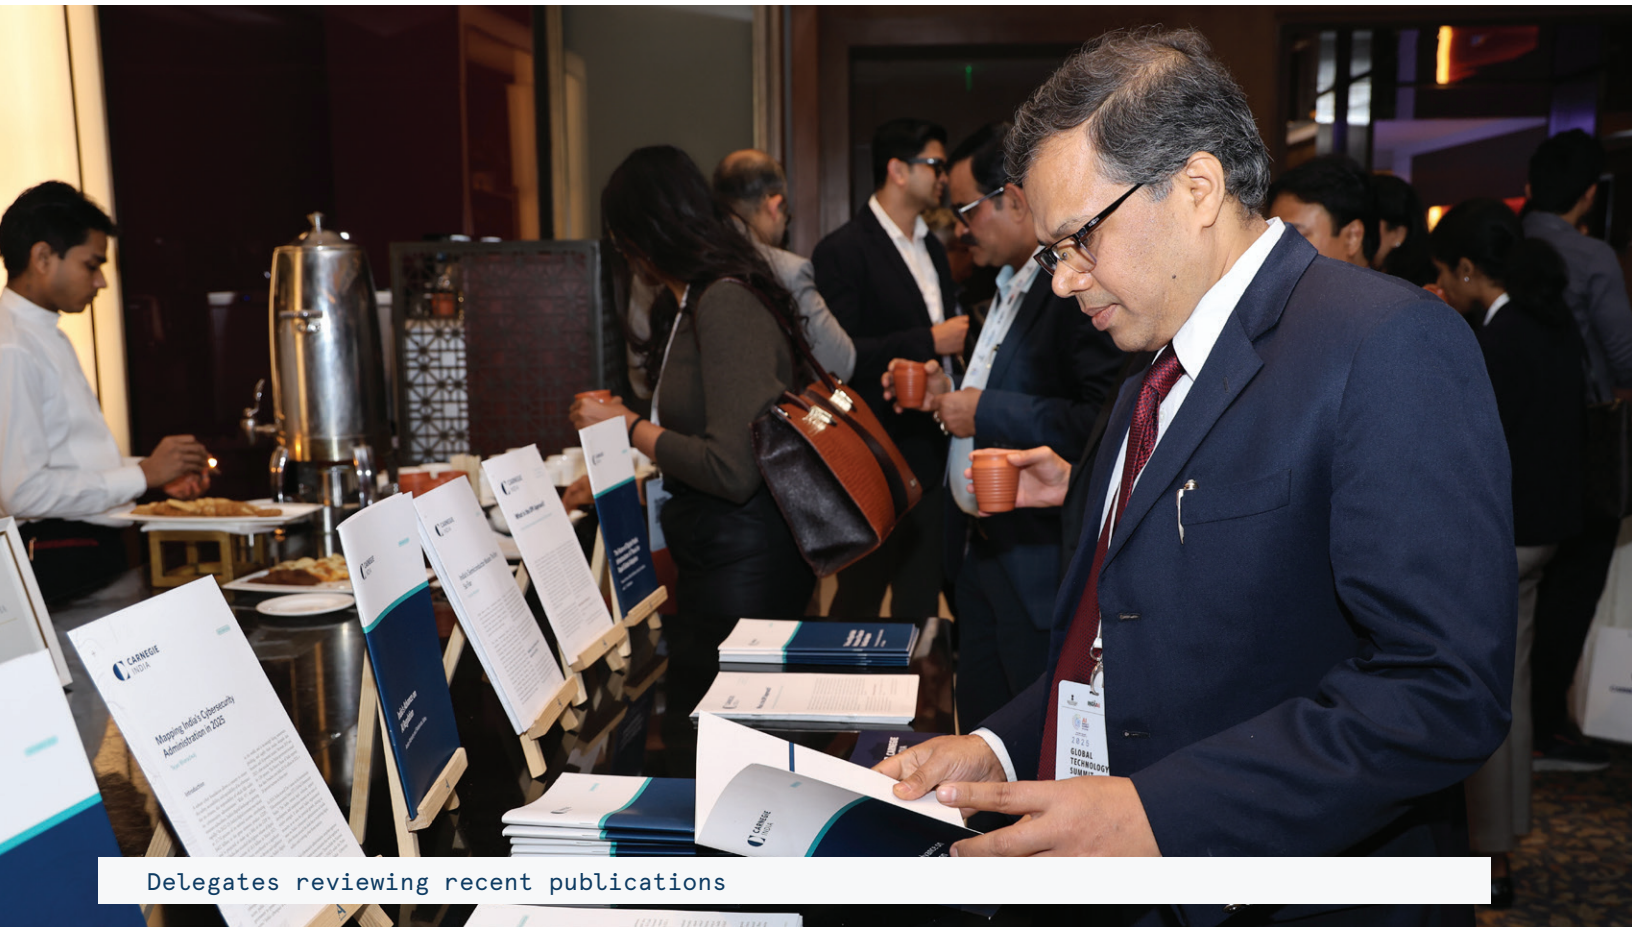

Delegates reviewing recent publications

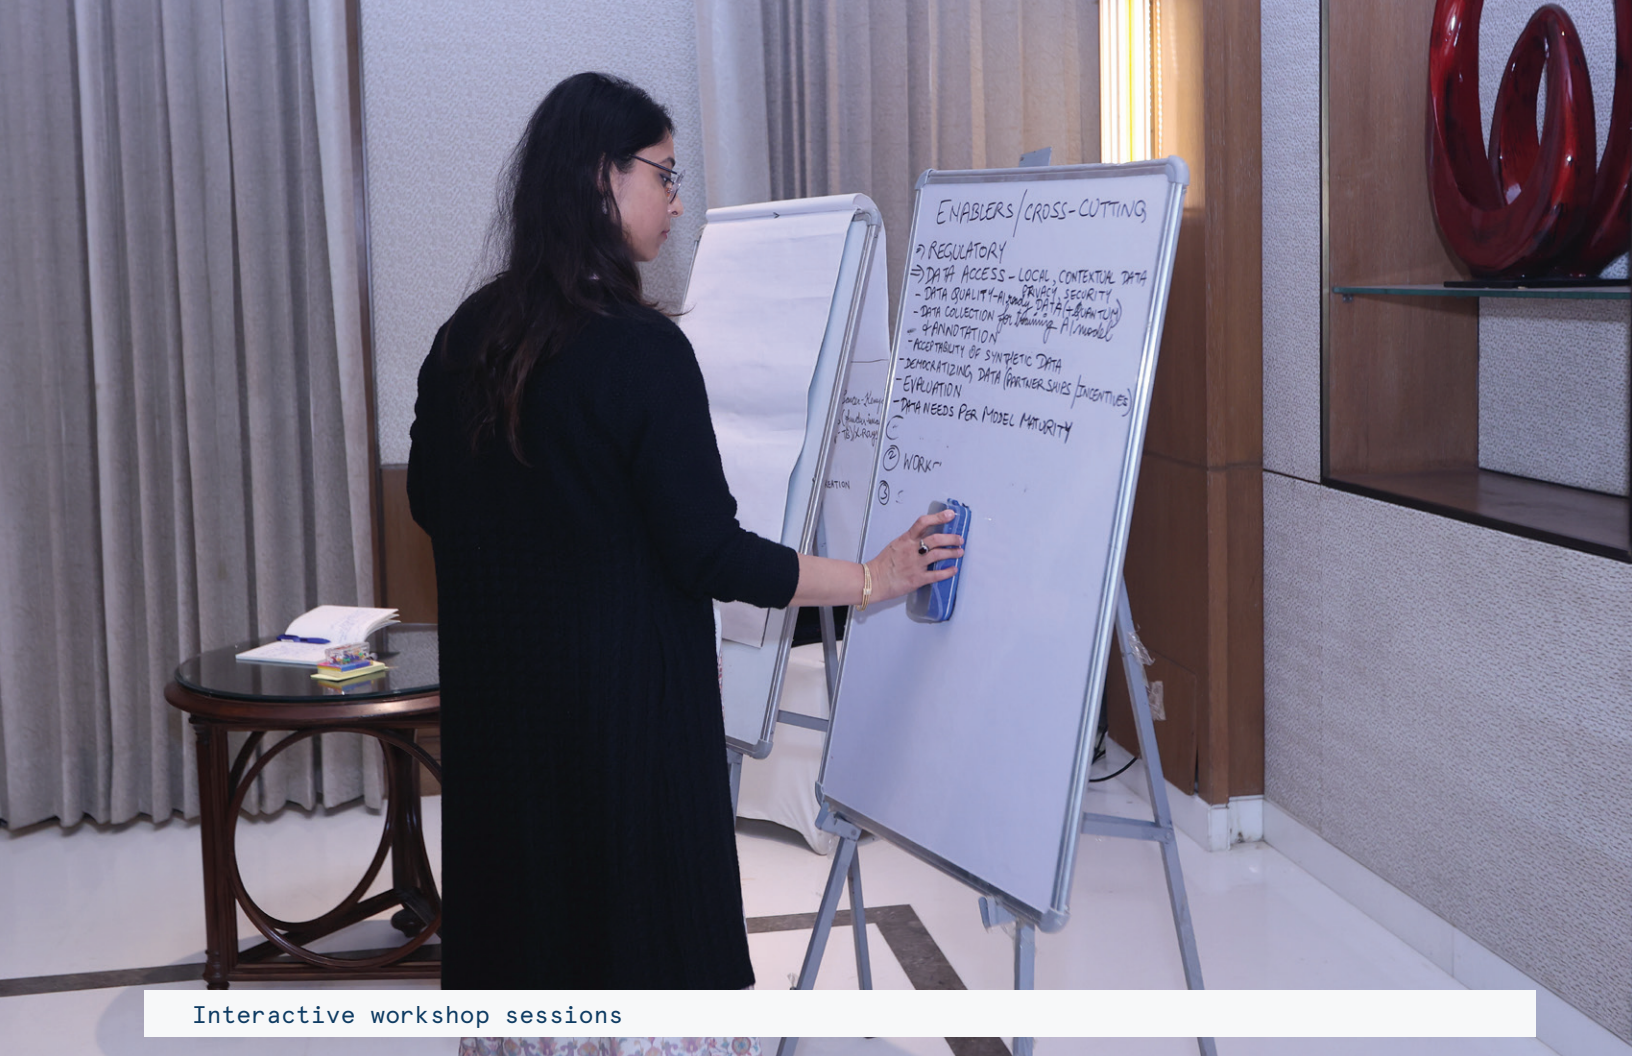

Interactive workshop sessions

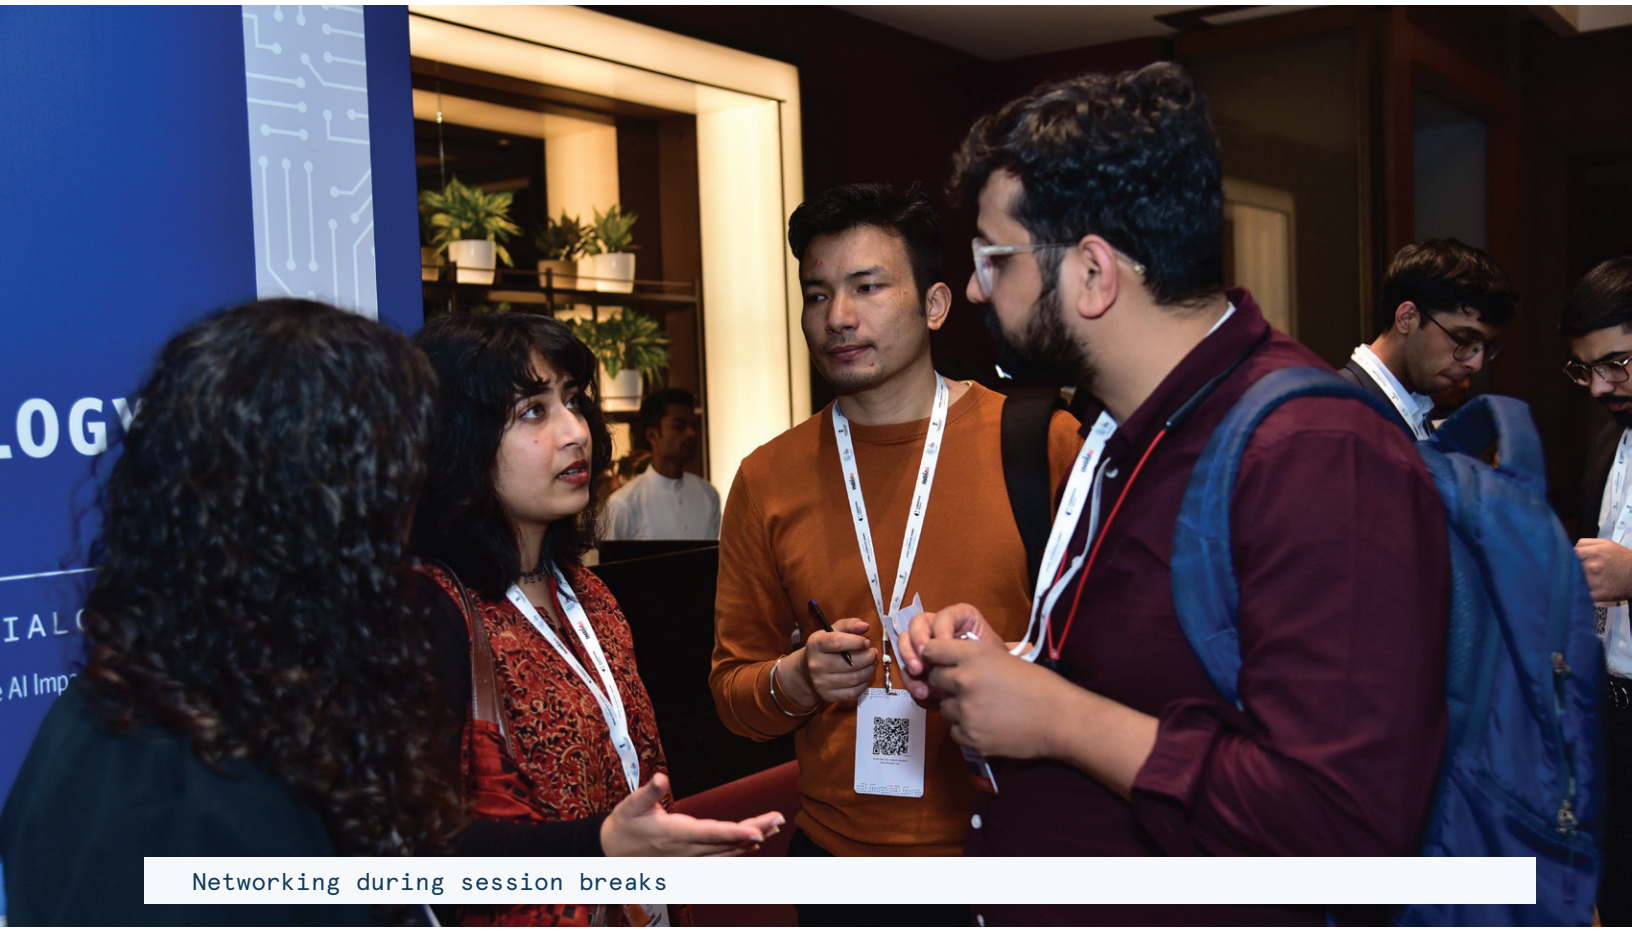

Networking during session breaks

AGENDA 2025

December 10, 2025

Opening Remarks

Speaker:

Aparna Ray, Joint Secretary (Policy Planning & Research Division), Ministry of External Affairs, Government of India

Address: Priorities of the AI Summit

Speaker:

Mohammed Y. Safirulla K., Director, IndiaAI Mission, Ministry of Electronics and Information Technology, Government of India (Virtual)

Closed-door Icebreaker Session

Closed-door Breakout Groups: Challenges and Opportunities in Scaling AI Use-Cases

- *AI in Healthcare*
- *AI in Agriculture*
- *AI in Education*

Closed-door Conversation: Take Stock, India AI Impact Summit

Closed-door Session: Use-Case (Dis)Content and Scale

Closed-door Session: U.S.-India Cooperation on AI Infrastructure, Ideas for Prosperity

December 11, 2025

Closed-door Workshop: Use-Case Framework for AI

Closed-door Parallel Session: Safe AI Use-Cases at Scale

Closed-door Parallel Session: AI for All, Catalysing a Global South Compute and Knowledge Commons

Closed-door Session: Open Reflections and Conclusion

Public Sessions: December 11, 2025

India AI Impact Summit Outcomes

Speakers:

***Debjani Ghosh**, Distinguished Fellow, NITI Aayog*

***Marisa Gerards**, Ambassador of the Kingdom of the Netherlands to India, Nepal and Bhutan*

***Saurabh Garg**, Secretary, Ministry of Statistics and Programme Implementation, Government of India*

Moderator:

***Archna Vyas**, Interim Country Director for the India Office, Gates Foundation*

Where AI Meets People: Impact Stories from the Global South

Speakers:

***Thiago Rached**, Co-founder & CEO, Letrus*

***El Mahdi Aboulmanadel**, Co-founder & CEO, DeepLeaf*

***Kidist Tesfaye**, Founder & CEO, YeneHealth*

***Benjamin Mwalimu**, Head of Technology, Jacaranda Health*

Moderator:

***Keyzom Ngodup Massally**, Director, AI Hub & Head of Digital & AI Programmes, UNDP*

Conversation with Sunil Wadhwani: AI for the Last Mile

Speaker:

***Sunil Wadhwani**, Wadhwani Institute for Artificial Intelligence; Founder, W Health Ventures*

Moderator:

***Shalini Kapoor**, Chief Strategist, Data and AI, EkStep Foundation*

The Spine that Powers Use-Cases: Compute and its Discontents

Speaker:

***Shikoh Gitau**, CEO, Qhala*

***Raymond U. Ononiwu**, Founder, Horus Labs*

***Raghu Dharmaraju**, CEO, ARTPARK at the Indian Institute of Science*

Moderator:

***Shankar Maruwada**, Co-founder and CEO, EkStep Foundation*

CO-HOST

Policy Planning & Research Division
Ministry of External Affairs
Government of India

PHILANTHROPY PARTNER

TATA TRUSTS

DIALOGUE PARTNER

EkStep

people + ai
An EkStep Initiative

KNOWLEDGE PARTNERS

OpenAI

Microsoft

amazon

Gates Foundation

salesforce

ANTHROP\C

LinkedIn

TATA CONSULTANCY SERVICES

CENTRAL SQUARE FOUNDATION

Meta

Policy Center for the Digital Economy

Qualcomm

THE BEST RUN SAP

HCLTech

IBM

ITU

OECD.AI Policy Observatory

AI Hub for Sustainable Development

UNDP

peak xv

MEDIA PARTNER

ANI
Asian News International

DIGITAL PARTNER

ThePrint
दिप्रिंट

OUTREACH PARTNER

asianetnews.com

CARNEGIE INDIA

सत्यमेव जयते
Policy Planning & Research Division
Ministry of External Affairs
Government of India

2025

GLOBAL TECHNOLOGY SUMMIT

INNOVATION DIALOGUE

AI IMPACT SUMMIT
भारत 2026 INDIA

सर्वजन हिताय | सर्वजन सुखाय
WELFARE FOR ALL | HAPPINESS OF ALL