AI KNOWLEDGE CONSORTIUM

DECODING AI: WHAT IT TAKES TO BUILD A FOUNDATION MODEL

ISSUE #1 | JUNE 2025

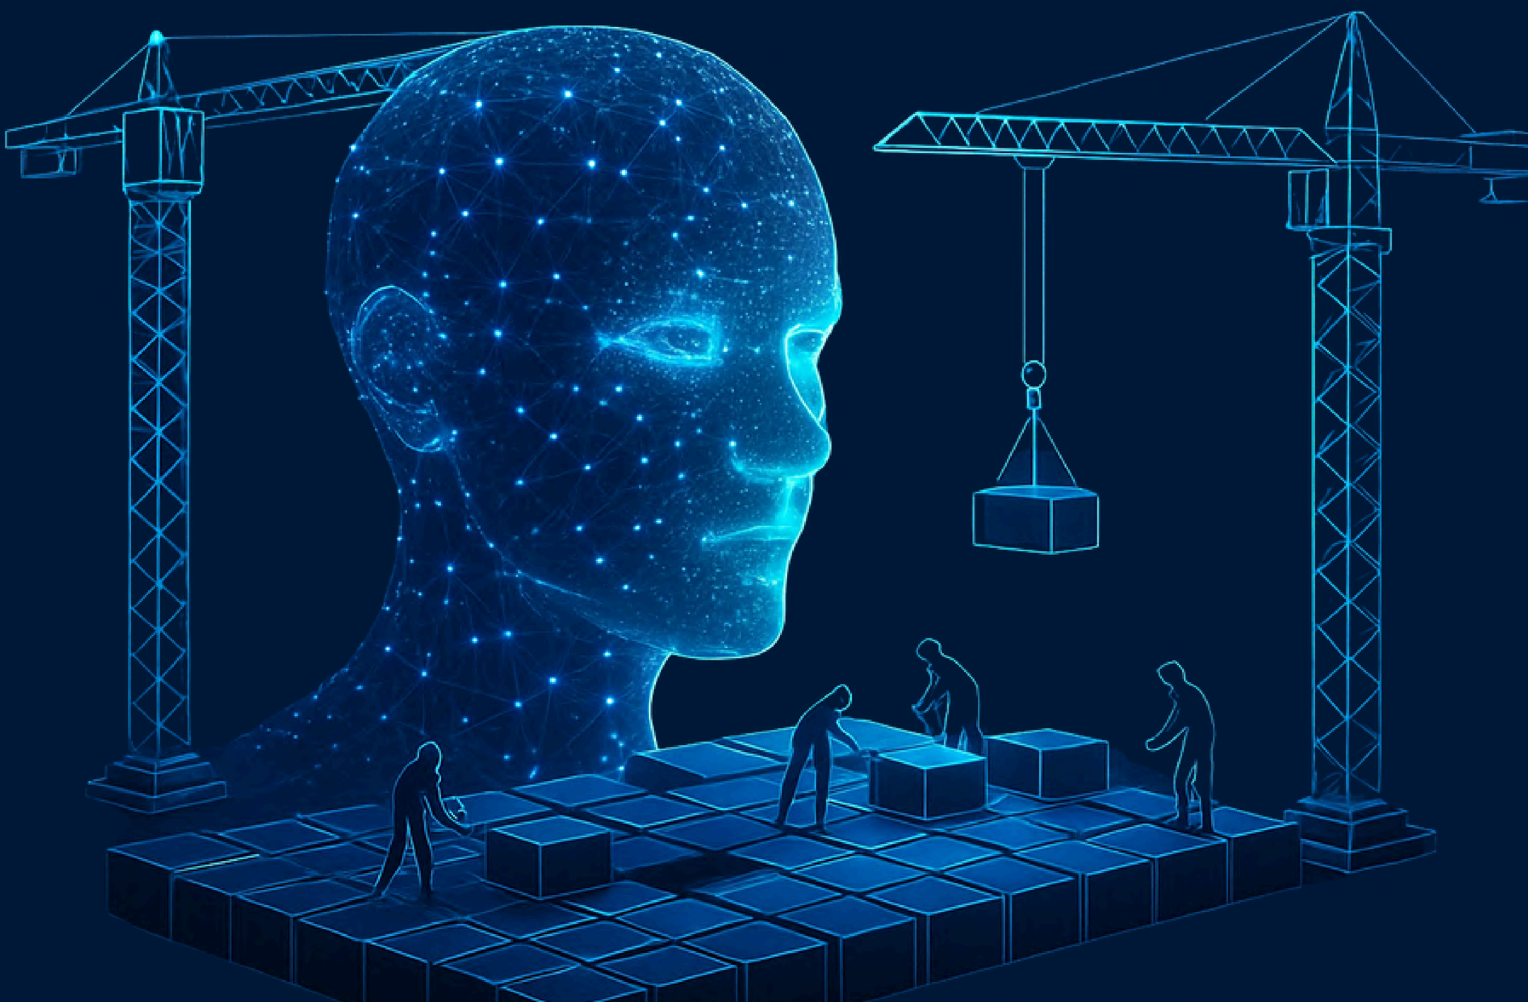

CONTEXT

Foundation models—large-scale AI systems trained on immense datasets—are at the heart of today's generative AI evolution. From the conversational capabilities of ChatGPT to cutting-edge vision-language tools, these models have the potential to transform every sphere, from healthcare and education to logistics and creative industries.

India stands at a pivotal moment in the global AI landscape. With growing policy momentum around building indigenous foundation models using Indian datasets, the question is no longer if but how. Meanwhile, rapid advances in countries like China have injected fresh urgency and competition into this space.

In this context, the AI Knowledge Consortium (AIKC) hosted a public discussion in New Delhi in April 2025, titled 'What it takes to build a Foundation Model.' The panel featured Navendu A, head of business at Krutrim; Saurabh Kumar Karn, founding member of Sarvam AI; Subodh Sharma, associate professor at IIT Delhi; and Pankaj Thakar, founder and investor at Pad Up Ventures. The session was attended by 100 stakeholders from academia, industry, and civil society organisations, among others. This report summarises the takeaways.

This session also marked the inaugural event of *Decoding AI*, the AI Knowledge Consortium's new public forum series. Designed to demystify complex concepts and foster broad engagement, Decoding AI will bring together experts, policymakers, industry leaders, and the public for open conversations on the technologies shaping our world. Through accessible discussions and interactive formats, the series aims to make AI more understandable, collaborative, and impactful for everyone in India.

INTRODUCTION

A foundation model is a large, general-purpose artificial intelligence (AI) system trained on diverse data and subsequently adapted to perform different tasks. Such a model can be used in many ways, including summarising documents, answering questions, or generating images and audio. Companies are racing to build the latest and most advanced foundation models, with several major developments in the past year. These include the release of updated models from established players like OpenAI (GPT-4), Google (Gemini 1.5), Meta (LLaMA 3), and Anthropic (Claude 3), entry of players like xAI with the Grok series, and DeepSeek making waves by claiming to achieve competitive performance while significantly reducing reliance on Graphics Processing Unit (GPU) heavy infrastructure.

These models are becoming critical infrastructure for economies and innovation systems. Consequently, they have evolved from a technological breakthrough to a matter of national priority worldwide. Over the past year, many countries have stepped up their efforts to build foundation models. Some have relied heavily on private firms, while others have pursued regulatory or state-led development:

United States: The US remains the global hub for foundation-model innovation, with most frontier models[1] having been developed by US companies. Executive orders by both the Biden and Trump administration[2] aim to strengthen US leadership in AI by revoking policies and directives that were viewed as barriers to innovation.[3]

China: China has made the development of foundation models a national priority, with President Xi Jinping emphasising the need for self-sufficiency in AI.[4] Over the past few years, it has rapidly scaled its foundation-model ecosystem. Several models, such as ERNIE 4.5 and ERNIE X1, were launched in 2025.[5] The DeepSeek R1 model, released in January 2025, caught the world's attention for achieving strong performance by focusing on memory optimisation, efficient training processes, and improved system design, rather than relying on increased GPU use - thereby reducing the reliance on imported technology.

European Union: The EU has adopted a regulation-first approach to foundation models. It enacted the world's first comprehensive AI law, which includes specific obligations for general-purpose AI (GPAI) models. The regulation follows a risk-based framework to ensure AI systems are safe, transparent, and accountable.[6] Although the EU hosts fewer frontier models, initiatives such as Aleph Alpha[7] in Germany reflect efforts to develop sovereign and transparent AI systems. Private efforts are also gaining momentum—new players like Mistral AI (France) aim to compete with OpenAI and DeepMind with their large-language model, Mistral Medium 3.[8]

India has been focusing on AI development through its IndiaAI Mission, launched in 2024 with a budget of ₹10,372 crore.[9] The mission is structured around seven key pillars including the IndiaAI Compute Capacity to develop scalable AI computing infrastructure, the IndiaAI Innovation Centre to develop indigenous foundational AI models and the IndiaAI Datasets platform to create a repository of high-quality, anonymised datasets for AI-model training and development.

SECTION I: FOUNDATION MODEL FUNDAMENTALS

Some Basic Concepts

A foundation model is a large, general-purpose AI system trained on diverse data that can later be adapted for different tasks. Such models can be employed in many ways, including summarising documents, answering questions, generating images, and producing audio.[12]

Foundation models are built to serve as a base layer. Rather than being narrow applications, they function as broad platforms on which other models and tools can be built. A company might train one foundational model and then fine-tune it for separate use cases, such as customer support, translation, or legal document analysis. For example, MedAlpaca is a model fine-tuned from Meta's LLaMA model. It was trained specifically on medical question-answering data using instruction tuning, thereby transforming a general-purpose foundational model into a healthcare-specific system.[13]

Tokens are the units of input the model processes. They are often smaller than full words. The model does not learn from whole sentences but from statistical patterns among tokens. The total number of tokens used during training directly affects model strength. Tokenisation refers to splitting raw text into discrete tokens. The tokeniser's design crucially influences how a model processes different scripts, sentence structures, and languages. For Indian language models, improved tokenisation is essential to accommodate local linguistic diversity. Because Indian languages employ distinct scripts and complex word forms, standard tokenisers, which are mostly developed for English, often perform poorly.

Parameters are the adjustable internal weights that the model uses to decide which token should appear next.[14] Input data is tokenised and passed through these parameters, which together generate the model's output. The more parameters a model contains, the more complex relationships it can capture. GPT-3 had about 175 billion parameters,[15] whereas newer models are expected to exceed a trillion.

Training data quality is equally important. Foundation models are generally trained on vast, diverse datasets with the aim of capturing patterns across domains rather than within a single task. After this pre-training, a model can be fine-tuned—that is, adjusted slightly—for a specific application or domain. For Indian models, however, access to large-scale, high-quality data in local languages remains a major challenge, as most widely used datasets are dominated by English and western content.

LLMs, SLMs and their use cases

At the most basic level, large language models (LLMs) are deep-learning systems trained to predict the next word in a sequence. Such models are often described as "sequence-to-sequence" because they process part of the input and then generate the next segment of output. As they are trained on more data, they learn statistical patterns; for example, which words or phrases tend to co-occur and in what order.

An LLM is trained on vast quantities of textual data and contains millions or even trillions of parameters. The breadth of training data and the sheer parameter count allow these models to learn abstract patterns across diverse text sources, making them broadly general-purpose.

While many foundational models are LLMs today, not all LLMs are foundational. Many are fine-tuned versions of existing models with narrow use cases. In general, LLMs are trained on large-scale text corpora to perform language-related tasks. Foundation models are trained on diverse, multimodal data (e.g., language, audio, and visual content) to support a wide range of downstream applications.[16] By contrast, small language models (SLMs) have far fewer parameters (usually one to two billion) and therefore require fewer resources to train and deploy. They are typically more task-specific and are often used for narrow applications such as document tagging or single-language translation. **SLMs are quicker to deploy and easier to fine-tune, making them a practical starting point for many organisations. However, they cannot match the breadth or generalisation ability of LLMs.**

Many in India are focusing on developing SLMs because the country faces resource constraints, training LLMs - which demand extensive computational power delivered by GPUs - is prohibitively expensive. India's unique linguistic and deployment needs also bring challenges. **With 22 official languages and significant regional variation, developing a single universal model is neither feasible nor necessary in the short term. SLMs instead allow for targeted language support and domain-specific fine-tuning, which is ideal for use cases in agriculture, education, law, and public-service delivery**. The domestic industry has concentrated on building and deploying smaller models that are locally useful and easier to train. Over time, these efforts are expected to evolve towards larger, more general-purpose models.

Does India Need a Foundation Model?

India's desire to build a domestic foundation model, evident in its policy push, it is also a response to structural and strategic risks and opportunities. Many stakeholders believe that exclusive reliance on global models—whether proprietary or open-source—limits India's ability to shape development trajectories. It exposes the country to licensing changes, output restrictions, and dependencies on companies or nations whose priorities differ from its own.

A locally built foundation model can be aligned with India's needs and values. It can be trained to support Indian languages, local use cases, and public-sector applications that international models may overlook. It would also give Indian institutions greater control over safety, alignment, and regulatory responses—concerns that are becoming more acute as diffusion policies and export controls grow increasingly unpredictable.

Most state-of-the-art (SOTA) models today are trained on 200–300 trillion tokens and contain 10–20 trillion parameters. These figures are far beyond current training capacities within the country. **Most stakeholders agree that while India can aspire to reach this scale in the long term, the immediate priority should be to build models that balance performance, cost, and local relevance.**

Questions for Future Research

1. How do SLMs developed in India compare with LLMs, when fine-tuned for specific use-cases, in terms of performance, training cost, and deployment feasibility?
2. What categories of risk—such as access restrictions, or licensing instability—might arise from relying on foreign foundational models, and how do these compare with the costs and risks of building India's own?

SECTION II: BUILDING BLOCKS FOR A FOUNDATION MODEL

Compute Infrastructure

Developing a foundation AI model requires substantial computational resources, particularly GPUs. GPUs are specialised hardware designed for parallel processing, enabling them to handle multiple operations simultaneously. This capability is essential for training large AI models, which involve processing vast amounts of data and complex computations. For example, training GPT-3 and GPT-4 required at least 20,000 GPUs—with the processing abilities of Nvidia's H100 series.

What matters even more than the raw GPU count is how these GPUs are arranged and accessed. Modern foundation models are trained on GPU clusters, which are large-scale, interconnected banks of GPUs optimised to work in parallel. In India, the government has announced plans to procure 10,000 GPUs under the IndiaAI Mission to bolster the country's AI infrastructure.[17] However, only about 2,000 of these have so far been organised into clusters, which are essential for efficient, large-scale AI training.

Private enterprises are also investing heavily. Ola Krutrim, for instance, has committed ₹10,000 crore to develop AI computing clusters, including systems such as GB200 and NVL72, leveraging NVIDIA technology.[18] **Traditionally, most organisations access compute through hyperscalers, which include cloud giants such as AWS, Google Cloud, and Azure that offer AI infrastructure as a service, including ready access to GPUs, model-training environments, and scaling tools.** According to Krutrim, its approach offers it greater control over resources, cost efficiencies, and the ability to tailor infrastructure to specific AI workloads.

Network Architecture

Network architecture can be just as important as the number of GPUs. Network-architecture optimisation refers to how different nodes, such as GPUs or compute servers, are connected in a training setup. In large-scale AI training, at least 89–90 percent of the nodes must be operational at all times for efficiency. In AI infrastructure, this includes designing GPU networks to minimise latency, and avoiding bottlenecks.

Simply providing high-performance computing resources does not guarantee their optimal utilisation. India's HPC systems have historically suffered from underutilisation; usage statistics show utilisation rates hovering around 70 percent for most systems.[19] Many of the tasks run by researchers, students, or institutions are not compute-intensive enough to exploit the systems' capacity. Consequently, valuable infrastructure is tied up with low-level workloads that do not justify their use. As India builds AI-focused GPU infrastructure, a similar pattern could emerge: availability without meaningful utilisation. To prevent this, quality-control mechanisms should be introduced when granting access, particularly in shared or virtualised systems.

India-Specific Data Sets

A foundation model is only as good as the data it is trained on. One of the biggest challenges India faces is the lack of high-quality data in Indian languages. Globally, datasets like Common Crawl which is an open archive of billions of webpages scraped from the internet, are widely used to train large language models. But these collections are dominated by English and other western languages. So, they are insufficient for India's context.

While initiatives such as AI4Bharat and AIKosh are working to curate Indian language datasets, there are several structural challenges. Data collection and annotation are expensive and resource-intensive. Much of India's linguistic data is not digitised in the form of archives, printed records, or unstructured speech content. Some of it is also locked in user generated and curated digital platforms, where accessing data for AI training raises intellectual property concerns.

Building accurate linguistic models requires more than just the volume of data. Diversity across dialects, regional variations, and social contexts is essential. Metadata—such as speaker background, geography, and usage context—is often missing. **Designing effective tokenisers, which convert text into machine-readable units, is especially complex for Indian scripts and requires a systems-level engineering approach. Tokenisation errors can severely reduce model performance in languages with rich morphology and flexible word order.**

To address these challenges, India must invest in creating open-source data sets. Open access enables external validation, improves the quality of data sets through feedback loops, and builds credibility for the organisations leading such efforts. **Consortiums of companies, government, and civil society should pool resources to create diverse, high-quality, open-source data sets to power AI in India. This will also accelerate research in several domains, such as science, health, and education.**[20]

Questions for Future Research

1. What are the ways to democratise systems design capabilities in India's AI ecosystem, including network architecture of public compute clusters?
2. What are the institutional and legal barriers to aggregating Indian-language data from fragmented sources and what governance models could streamline ethical access and reuse for AI training?

SECTION III: BEARING THE COSTS

Building a foundation model with trillions of parameters requires substantial upfront and recurring capital outlay. The costs are driven by compute facilities, multilingual complexity, emerging talent, and research development.

- Compute cost: Training a SOTA model, as previously noted, would require 200–300 trillion tokens and involve 10–20 trillion parameters. The capital cost of procuring and maintaining this infrastructure can exceed several hundred million dollars. As a reference, OpenAI's GPT-4 cost over $100 million to train[21] and Grok 3 cost between $3 to 4 billion.[22] The costs will escalate if there is a need for custom tokenisers (as might be needed for Indian multilingual models) or if network architecture isn't optimised. India currently has access to approximately 10,000 GPUs, of which only about 2,000 are configured in clusters suitable for training large models.[23]

- Data cost: Data-acquisition and curation costs in India are both unique and complex. Most Indian-language content is either not digitised, or exists only in audio or video formats. Creating structured, clean, and diverse datasets in Indian languages significantly increases the cost of data collection and augmentation. Licensing rights, copyright ownership, and fair-use policies also affect data availability and costs.

- Skill development cost: India's current talent pool is concentrated largely in application-layer roles. Building a foundation model will require both skill-building and the development of a broader talent pool; these, in turn, must be supported by academic courses aligned with industry demand.

Building a foundation model is expensive, and the ecosystem is too layered for any one stakeholder to bear the cost alone. Most stakeholders agree that only a balance of public investment, private capital can make a foundation model viable in the Indian context.

The Role of the Government

The government must play a pivotal role in laying down the public infrastructure by accelerating GPU procurement and establishing clusters for training large models. Through the IndiaAI Mission, the government has already planned a national compute cluster of 18,693 GPUs for immediate use by start-ups and researchers via the AI Compute Portal.[24] Nearly half of the ₹10,300 crore budget allocated to the Mission is earmarked for this facility. Additionally, public funding can support digitising archives, standardising language data, and creating open data resources (like Common Crawl and AIKosh). The government can coordinate by preventing duplicate data collection, ensuring data quality, and supporting shared infrastructure.

Academic institutions and sandbox organisations must receive subsidised access to the AI foundation model to enable testing and innovation.[25] The government should also consider tax waivers for companies investing in AI infrastructure.

In addition to investing in computing infrastructure, government departments and agencies should also drive the research and development of AI systems for specific purposes / end-uses. The US Defense Advanced Research Projects Agency runs the AI Next and AI Forward campaigns to advance AI research and development for national security.[26] India's Department of Science and Technology and the newly constituted Anusandhan National Research Foundation are working to replicate and adapt this model to the Indian context. They recently announced [27] the following initiatives:

- AI for Science Mission: use machine learning to accelerate discoveries in physics, chemistry, and biology.
- India AI Open Stack: develop a foundational AI architecture, embedded with scientific and engineering models, to support Indian research.

Industry collaborations, such as T-Hub and the Department of Science and Technology's Machine Learning and Artificial Intelligence Technology Hub which are supported by the Department of Science and Technology and the government of Telangana, should also be promoted. Such initiatives foster AI innovation, create job opportunities, and provide sandbox environments for start-ups.[28]

Finally, the government should support and incentivise public-private research collaborations, offer higher fellowships to PhD candidates, and streamline intellectual property-sharing frameworks to promote AI systems research.

The Role of Private Institutions

In the US and China, government support has been complemented by substantial private-sector investment from companies such as Microsoft and High-Flyer.[29] Although India is beginning to take similar steps, with Sarvam AI and Ola Krutrim showing early interest,[30] **Venture capitalists do not appear to favour investing in a foundation model. This reluctance is largely because such models are expensive to build and their monetisation is uncertain.** Even OpenAI, the developer of ChatGPT, has acknowledged the financial strain of running its foundational models despite their widespread use and paid subscriptions.[31]

Venture capitalists are therefore concentrating on applications and services built on top of existing models that can be deployed quickly and offer a faster return on investment. Large corporations, particularly those that require proprietary AI models should nevertheless consider investing in foundation-model development.

Improve Cost Effectiveness

Companies developing AI models should think about lowering costs by optimising existing technologies and look at global best practices. Consider the example of AI chips: the extent to which an AI system maximises output is called yield efficiency. Higher yield efficiency, combined with minimal resource consumption, reduces the cost per functioning AI chip. Taiwan Semiconductor Manufacturing Company (TSMC) and Samsung both manufacture 3-5 nm chips; however, TSMC's yields exceed 85 percent, compared with Samsung's 15 percent.[32] India lacks domestic high-end semiconductor fabrication. Indian model makers should collaborate with manufacturers such as TSMC and design chips specifically for AI workloads, thereby reducing per-unit training costs. Costs can be lowered further by prioritising the country's latest semiconductor-manufacturing initiative to produce indigenous chips.[33]

DeepSeek reduced costs by optimising across multiple levels. For instance, it employed effective memory management, which reduced the compute load during training. It also streamlined hardware usage and adjusted training schedules to make more efficient use of available resources. Through architectural and training innovations, **it is possible to attain near-SOTA performance at a fraction of traditional costs. India can build large-scale capability with limited infrastructure if the right areas are optimised carefully.**

Questions for Future Research

1. How can India leverage its strengths in application development to help commercialise as well as indigenise other layers of the AI stack, from chips to the cloud?
2. What are the government's priority use cases for AI that can justify public investment and enable cost recovery?

SECTION IV: COLLABORATIVE CAPACITY BUILDING

Academia-Industry-Government Linkages

Stakeholders agree that since AI is evolving rapidly, academic institutions often lack the time to design, build, and then teach new courses and there is a shortage of high-quality courses and infrastructure that allow students to experiment and gain hands-on experience with AI technology. There are around 670 AI and data-science engineering degrees in the country.[34] However, the practical value of these programmes, and the extent to which they expose students to technical problem-solving, remain doubtful. Even leading institutions such as the IITs face resource constraints when procuring the infrastructure required for students to experiment and gain first-hand experience with foundational models.

Existing academic incentive structures are inadequate to encourage innovation and experimentation among research scholars. Insufficient funding, combined with limited family support, accommodation, and other facilities, discourages many from remaining in research.

Academia must collaborate with both industry and government. Creating Centres of Excellence (CoEs) within universities—where industry and government work together—can furnish students with the necessary infrastructure and supplement courses. India's 2025 budget envisages the establishment of a CoE in artificial intelligence for education, indicating that such collaborations are already under way. These models have been successful elsewhere: for example, China has played an active role in developing a skilled AI talent pool by providing well-funded research grants and scholarships, and by establishing linkages between academia and industry.[35]

Stronger academia–industry linkages are needed to create continuous feedback loops that allow academic researchers to test their models in real time, access sufficient training data, and ensure that those models perform well in practical settings. Some initiatives such as Krutrim's AI Lab have been kickstarted in this direction.

The Talent Conundrum

Most agree that India currently possesses adequate talent in AI applications. Nevertheless, the country requires more professionals skilled in systems engineering, design engineering, low-level code, and GPUs. Specialists in statistics, probability theory, logic, and linear algebra will also be essential for designing foundation models.

Estimates indicate that India's AI sector could have 2.3 million job openings by 2027, yet only 1.2 million persons with the requisite talent, presenting an opportunity to reskill and upskill more than 1.1 million workers.[36] A key source of talent could be the fungible talent pool in the IT industry that circulates within a few companies in search of incremental pay rises, and benefits. These workers can be reskilled for design- and systems-level engineering roles.

Questions for Future Research

1. How can academia and industry collaborate to ensure that academic curricula keeps pace with the development of AI technology?

ENDNOTES

1 A frontier model may be defined as a highly-capable general-purpose model that can perform a wide range of tasks and it matches or exceeds the capabilities of the most advanced models that currently exist, Frontier AI: capabilities and risks – discussion paper - GOV.UK, https://www.gov.uk/government/publications/frontier-ai-capabilities-and-risks-discussion-paper/frontier-ai-capabilities-and-risks-discussion-paper#fn:2

2 Former US president Joe Biden signed Executive Order 14110, titled “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”, in October 2023. The current Trump administration issued Executive Order 14179, titled “Removing Barriers to American Leadership in Artificial Intelligence”, in January 2025.

3 Executive Order 14179, titled “Removing Barriers to American Leadership in Artificial Intelligence”, in January 2025

4 James Pomfret and Summer Zhen, ‘China’s Xi calls for self sufficiency in AI development amid U.S. rivalry’, Reuters, (2025), https://www.reuters.com/world/china/chinas-xi-calls-self-sufficiency-ai-development-amid-us-rivalry-2025-04-26/

5 ‘Baidu rolls out Ernie 4.5 and Ernie X1 AI models, promises advanced reasoning and cost efficiency’, Mint, (2025) https://www.livemint.com/ai/artificial-intelligence/baidu-rolls-out-ernie-4-5-and-ernie-x1-ai-models-promises-advanced-reasoning-and-cost-efficiency-11742207126503.html

6 Implementation Timeline, EU Artificial Intelligence Act (2024), https://artificialintelligenceact.eu/implementation-timeline/

7 Lighthouse for sovereign AI development in Germany: Aleph Alpha strengthens collaboration with the IPAI ecosystem for open and application-oriented research and development in GenAI’ , Aleph Alpha, (2024), https://aleph-alpha.com/lighthouse-for-sovereign-ai-development-in-germany-aleph-alpha-strengthens-collaboration-with-the-ipai-ecosystem-for-open-and-application-oriented-research-and-development-in-genai/

8 ‘Mistral announces new AI model Medium 3 at 8x lower cost’, The Indian Express (2025), https://indianexpress.com/article/technology/artificial-intelligence/mistral-launches-medium-3-ai-model-at-8x-lower-cost-9991190/

9 Press Information Bureau, ‘Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem’, 2024, https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012357

10 Soumyarendra Barik, ‘Centre selects start-up Sarvam to build country’s first homegrown AI model’, Indian Express (2025), https://indianexpress.com/article/business/centre-selects-start-up-sarvam-to-build-country-first-homegrown-ai-model-9967243/

11 Framework for Artificial Intelligence Diffusion, Federal Register, The daily Journal of the United States Government (2025), https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion

12 FAQ’s on foundational models and generative AI, The White House (2023), https://www.ai.gov/wp-content/uploads/2023/09/FAQs-on-Foundation-Models-and-Generative-AI.pdf

13 Chaoyi Wu and others, ‘PMC-LLaMA: Towards Building Open-source Language Models for Medicine’, Cornell University (2023), https://arxiv.org/abs/2304.14454

14 Olivia Shone, ‘Explore AI models: Key differences between small language models and large language models’, Microsoft (2024), https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/11/11/explore-ai-models-key-differences-between-small-language-models-and-large-language-models/

15 OpenAI Presents GPT-3, a 175 Billion Parameters Language Model, NVIDIA Developer Blog (2020), https://developer.nvidia.com/blog/openai-presents-gpt-3-a-175-billion-parameters-language-model/

16 Helen Toner, ‘What Are Generative AI. Large Language Models, and Foundation Models’, CSET (2023), https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models/

17 Press Information Bureau, ‘Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem’ (2024), https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012357

18 Nabodita Ganguly, ‘Ola Krutrim to Invest Rs 10,000 Crore in Krutrim AI Over Next One Year’, Outlook Start-Up (2025), https://www.outlookbusiness.com/start-up/news/ola-krutrim-to-invest-rs-10000-crore-in-krutrim-ai-over-next-one-year

19 Annual Report, National Param Supercomputing Systems (2024), NSM-AR-book-design-02425_updated_11zon.pdf

20 For instance, “The Well” is a 15-terabyte physics-simulation data set that can advance deep learning in physics. https://polymathic-ai.org/the_well/

21 Will Knight, ‘OpenAI’s CEO Says Age of Giant AI Models Is Already Over’, WIRED (2023), https://www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over/

22 Elon Musk Hyping Up xAI's Grok 3 with Massive GPU Investment', Open Data Science, Medium (2024), https://odsc.medium.com/elon-musk-hyping-up-xais-grok-3-with-massive-gpu-investment-45b553436f47

23 Press Information Bureau, 'India's AI Revolution' (2025), https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810

24 Press Information Bureau, 'India's AI Revolution' (2025), https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810

25 Press Information Bureau, 'India's AI Revolution' (2025), https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810

26 'Remaining the future of artificial intelligence for national security', DARPA (2022), https://www.darpa.mil/research/programs/ai-forward

27 Press Information Bureau, 'Dr. Jitendra Singh Reviews DST Projects, Calls for AI led innovations, Support for Deep-Tech Startups' (2025), https://www.pib.gov.in/PressReleasePage.aspx?PRID=2127062

28 Kajal Mehra, 'T-Hub and DST inaugurate AI/ML CoE, spearheading India's tech progress', timestech.in (2024), https://timestech.in/t-hub-and-dst-inaugurate-ai-ml-coe-spearheading-indias-tech-progress/

29 'ChatGPT maker OpenAI closes largest private tech deal of all time with $ 40 billion funding at $300 billion valuation and a clause' , Times of India (2025), https://timesofindia.indiatimes.com/technology/tech-news/chatgpt-maker-openai-closes-largest-private-tech-deal-of-all-time-with-40-billion-funding-at-300-billion-valuation-and-a-clause/articleshow/119860673.cms;
Amy Hawkins, 'Who is behind DeepSeek and how did it achieve its AI 'Sputnik moment'?', The Guardian (2025), https://www.theguardian.com/technology/2025/jan/28/who-is-behind-deepseek-and-how-did-it-achieve-its-ai-sputnik-moment

30 'What are foundational models, examples of common ones?', Business Standard (2025), https://www.business-standard.com/technology/tech-news/explained-what-are-foundational-models-examples-of-common-ones-125012301210_1.html;
'Sarvam AI to Build India's first homegrown AI foundational model, announces Ashwini Vaishnaw', Moneycontrol (2025), https://www.moneycontrol.com/news/business/sarvam-ai-to-build-india-s-first-homegrown-ai-foundational-model-announces-ashwini-vaishnaw-13005616.html

31 Tobias Mann, 'Even at $200/mo, Altman admits ChatGPT Pro struggles to turn a profit', Hewlett Packard Enterprises (2025), https://www.theregister.com/2025/01/06/altman_gpt_profits/

32 'TSMC to dominate 2nm sector', Taipei Times (2025), https://www.taipeitimes.com/News/biz/archives/2025/03/24/2003833926

33 'India's first homegrown semiconductor chip to launch by end of 2025: Ashwini Vaishnaw', India Today (2025), https://www.indiatoday.in/technology/news/story/indias-first-homegrown-semiconductor-chip-to-launch-by-end-2025-ashwini-vaishnaw-2706589-2025-04-09

34 Prof. Anindya Chakrabarti and others, 'Labour-force perception about AI: A study on Indian white-collar workers', Brij Disa Centre for Data Science and Artificial Intelligence, IIM-Ahmedabad (2024), https://www.iima.ac.in/sites/default/files/2024-09/Labour-force%20perception%20about%20AI%20Final%20version%20%281%29.pdf

35 Fan Wang and João da Silva, 'DeepSeek: How China's 'AI heroes' overcame US curbs to stun Silicon Valley', BBC (2025), https://www.bbc.com/news/articles/czepw096wy4o

36 Bain & Company, 'AI: The Ambitions Are Bold, but the Talent Is Scarce' (2025), https://www.bain.com/about/media-center/press-releases/20252/widening-talent-gap-threatens-executives-ai-ambitions--bain--company/

Secretariat

KOAN

Koan Advisory Group, a New Delhi-based public policy consulting firm provides secretarial support to the AIKC

Contact Us

For inquiries, partnerships, or to learn more about our work with the AIKC, please write to:
Secretariat@aiknowledgeconsortium.com

Address: B40, Soami Nagar, New Delhi, 110017