When Xi Jinping sits down with Donald Trump in Washington on September 24, artificial intelligence will be near the centre of the table. Washington has plenty to discuss: AI-enabled cyberattacks, semiconductor controls and allegations that Chinese laboratories have used American models to accelerate their own development.

But less than two weeks earlier in New Delhi, Xi presented a very different AI agenda. China proposed an open-source AI community for BRICS, a digital-industry cloud platform and an alliance to train engineers across the bloc.

The contrast reveals something larger than another chapter in the technological rivalry between China and the United States.

The geography of the AI race may be changing.

America retains formidable advantages at the technological frontier. Its companies build many of the world’s most capable models, design its most important AI chips and operate enormous cloud platforms. But the next stage of artificial intelligence will not be determined solely by who produces the most powerful model.

It will also be determined by who can make AI cheap enough, adaptable enough and accessible enough for billions of people to use.

On that measure, the emerging contest may carry an increasingly significant Asia Advantage — and China and India, for very different reasons, stand to gain the most.

For Washington, the AI debate increasingly revolves around technological leadership: who owns the best models, who possesses the fastest chips, who copied whose capabilities and how far ahead the United States can remain.

For much of Asia, Africa and Latin America, the questions are considerably more practical.

Can a health ministry in Jakarta download a model and run it on its own infrastructure? Can an Indian state adapt one across several local languages? Can an African university retrain it without paying dollar-denominated fees every time a student makes a query? Can governments keep critical AI systems operating if their relationship with Washington or an American technology company deteriorates?

These are questions of affordability, but increasingly they are also questions of sovereignty.

And Asia happens to contain the two countries best placed to answer them at scale.

China’s advantage is increasingly obvious.

Alibaba’s Qwen family has emerged as one of the world’s most widely adopted model ecosystems. Chinese models have also gained significant usage among American developers through platforms such as OpenRouter. Developers choosing Qwen, DeepSeek or Kimi are not necessarily making geopolitical statements. They are responding to familiar incentives: capable models, permissive licences and lower costs.

China can combine those models with something few countries possess simultaneously: manufacturing capacity, cloud infrastructure, enormous domestic demand, engineering talent and a government willing to treat technological ecosystems as instruments of international engagement.

The BRICS proposals extend that logic.

A country that trains its engineers on Chinese models does not simply acquire software. It acquires technical habits, evaluation methods, developer tools and institutional relationships.

Those connections can survive longer than any individual generation of AI.

Pedagogy is cheaper to export than silicon. It is also considerably harder to sanction.

Yet focusing exclusively on China’s rise misses the other half of Asia’s emerging advantage.

India occupies a remarkably different position.

It does not need to reproduce China’s model to benefit from the same transformation. India’s strengths lie in the layer where artificial intelligence moves from laboratory capability into mass deployment: software engineering, services, multilingual adaptation, digital public infrastructure and an enormous domestic market capable of testing AI applications at population scale.

The country also confronts precisely the problems that the next generation of AI must solve.

India does not have one linguistic market but many. It must build systems capable of operating across languages, income levels and uneven digital infrastructure. It needs AI that works not merely for wealthy consumers with premium subscriptions but potentially for farmers, small businesses, schools, hospitals and government services.

If the first phase of generative AI rewarded the ability to spend billions training frontier models, the diffusion phase may reward something different: the ability to make increasingly capable models useful at very low cost.

That plays considerably closer to India’s traditional technological strengths.

There is also a geopolitical advantage.

India does not fit neatly inside either an American or Chinese technology bloc. It can cooperate with American semiconductor and technology companies while developing domestic AI capacity. It can participate in BRICS initiatives without accepting Chinese technological leadership. And it can advocate technological autonomy for developing economies without possessing China's political baggage.

That makes India potentially more than another market for somebody else's models.

It could become one of the world's largest laboratories for adapting them.

Southeast Asia already demonstrates how this model might work. Singapore's SEA-LION initiative has developed models around the linguistic diversity of the region, while Indonesia has pursued Sahabat-AI for locally relevant applications. Such projects can draw from open foundations rather than requiring every country to spend billions developing a frontier model from scratch.

That is where open-weight AI changes the economics.

A proprietary model generally reaches Jakarta, Lagos or Nairobi as a service. Someone abroad hosts it. Someone abroad prices it. Someone abroad determines the conditions under which it remains available.

An open-weight model can arrive as something closer to technological capital. It can be downloaded, adapted and incorporated into domestic systems.

Once the weights are on local infrastructure, they are much harder to take away.

There are important limits to this argument.

Open weights are not synonymous with open source. A genuinely open ecosystem requires considerably more than downloadable parameters, including transparency around training, tooling and data. Nor does downloading a model eliminate infrastructure dependence.

Serving AI to tens or hundreds of millions of people still requires accelerators, data centres, electricity, cooling and skilled engineers.

That constraint may actually reinforce the Asia Advantage.

China can compete in infrastructure and hardware. India can compete in software, engineering, localisation and low-cost deployment. Southeast Asian economies can specialise in particular pieces of the semiconductor, data-centre and application stack.

What begins as competition among Asian powers could therefore coexist with a broader shift of AI's centre of gravity toward the region.

India and China will certainly compete within it.

New Delhi has little reason to allow a China-centred AI architecture to become the default for BRICS. India's own model ambitions, digital infrastructure and strategic concerns point toward technological autonomy rather than dependence on Beijing.

That competition could itself prove productive.

For developing countries, the relevant choice need not remain binary: pay American companies indefinitely for proprietary intelligence or adopt a Chinese technological ecosystem wholesale. India has an opportunity to help create a third model based around interoperable systems, domestic adaptation and technological sovereignty.

This is why Washington risks misunderstanding the next phase of the AI race if its response remains dominated by theft, export controls and technological ceilings.

Anthropic's allegations that Chinese laboratories conducted large-scale unauthorised distillation deserve examination. AI-enabled cyberattacks deserve serious negotiations. Advanced semiconductors inevitably have national-security implications.

But none answers the question being asked by governments across the developing world:

How do we obtain AI capability of our own?

America's technology companies possess much of the world's leading AI research. Asia increasingly possesses something equally important: the environment in which AI must become affordable enough to reach enormous populations.

The distinction matters because technological revolutions are rarely won only by inventors.

They are also shaped by adapters, manufacturers and distributors.

The United States dominated important layers of the internet economy, but Asia became indispensable to the physical production of the digital world. Artificial intelligence could produce another division of labour — except this time Asian countries may occupy considerably more of the software stack as well.

China wants to supply models, infrastructure and training.

India can supply engineers, applications, localisation and population-scale deployment.

Southeast Asia offers manufacturing networks, rapidly growing digital economies and

hundreds of millions of users.

Together, these constitute something that is still underestimated in Western discussions of the AI race: a vast continental market in which the overriding requirement is not necessarily the world's most intelligent model, but the world's most economical useful intelligence.

That is the Asia Advantage.

It does not guarantee Asian technological dominance. The United States retains enormous advantages in capital, research, chips and frontier models. China faces questions about trust and technological dependence. India still lacks sufficient domestic compute and must translate its enormous talent base into indigenous AI intellectual property and infrastructure.

But the economics of AI are moving rapidly.

Models are becoming cheaper. Open weights are spreading. Distillation is reducing the cost of reproducing capabilities. Smaller models are becoming more useful. And the competitive battleground is gradually moving from simply creating intelligence to deploying intelligence.

That transition favours countries that know how to operate at scale and under cost constraints.

Few countries understand scale better than China.

Few understand low-cost software deployment better than India.

By the time China assumes the BRICS chairship in 2027, the success of Beijing's New Delhi proposals should therefore be measured in tangible outcomes: models downloaded, engineers trained, compute made available and government systems actually deployed.

India’s progress should be measured by an equally demanding standard: how much of this emerging ecosystem it can localise, build and ultimately export rather than merely consume.

The Trump-Xi summit will understandably focus on the contest between the world’s two largest powers. But beyond Washington, another contest is already taking shape.

It is a competition to supply billions of people with usable artificial intelligence.

China enters that race with infrastructure and increasingly competitive models. India enters it with engineers, software capability, digital public infrastructure and extraordinary scale.

America may continue building many of the world’s most powerful AI systems.

But if the decisive question of the next decade becomes not who builds the smartest model? but who can put useful intelligence into the hands of the most people at the lowest cost?, the geography of advantage begins to look very different.

It begins to look Asian.