

No Seat at the Table: Why the Global South is Absent from the AI Governance Conversation

Nabib Bin Zahid¹



Introduction

In May 2023, more than 350 artificial intelligence researchers, engineers, and company executives put their names to a single sentence and released it to the public: mitigating the risk of extinction from AI should be a global priority alongside pandemics and nuclear war.² The signatories included the chief executives of OpenAI, Google DeepMind, and Anthropic — the same people who spend their working days building the systems that sentence warned about. A few months earlier, Geoffrey Hinton, one of the researchers whose work made the current wave of AI possible,

¹ Nabib Bin Zahid is a Research Assistant at the Bangladesh Institute of Peace and Security Studies (BIPSS). He has completed his B.S.S in International Relations from the University of Rajshahi.

²Center for AI Safety, "Statement on AI Risk," May 30, 2023, <https://www.safe.ai/statement-on-ai-risk>.

resigned his position at Google so that he could speak about these risks without implicating his employer.³

None of this happened in Dhaka, Nairobi, Jakarta, or São Paulo. It happened in San Francisco, in London, and inside a small cluster of laboratories that between them command more computing power than most national governments. That geography, not the extinction-risk headline it produced, is the actual subject of this piece. If the people who design frontier AI systems are themselves uncertain whether they can predict or fully control what they have built, then the roughly six billion people living outside that small circle of laboratories and governments are not merely absent from a policy conversation. They are exposed to a risk they did not define, governed, to the extent they are governed at all, by rules they had no part in writing.

South Asia's stake in this is proportionally larger than almost anywhere else, a fact that gets little attention in a debate dominated by superintelligence timelines and compute clusters. The region is home to close to two billion people, roughly a quarter of humanity, on a landmass a fraction the size of the countries currently writing AI's rules.⁴ Bangladesh alone has one of the highest population densities of any country its size in the world; a labor-displacing technology, a misinformation-amplifying one, or a financial-system error plays out differently at that density than it does in a country with a fraction of the people per square kilometer. Weak regulatory capacity means South Asian states are among the least equipped to absorb these effects and among the least present when the mechanisms meant to prevent them are designed.

The Global South's relationship to frontier AI is one of extraction without agency: it supplies the data, the labor, and the overwhelming majority of the user base on which these systems are built, yet holds no seat at the tables, from the UN's AI advisory bodies to the G7-driven safety summits, where their safety thresholds, liability rules, and epistemic defaults are decided. This exclusion is not imposed from outside alone. It is compounded by the Global South's own failure to organize a collective voice, even as the EU and individual middle powers have secured theirs. The cost is not merely political absence from governance, but epistemic absence from the systems themselves — whose language, whose context, and whose assumptions count as default, decided entirely by

³Cade Metz, "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead," New York Times, May 1, 2023.

⁴United Nations Department of Economic and Social Affairs, World Population Prospects 2024 (South Asia's population is approximately 2 billion, or roughly 25 percent of the world total).

others. And this asymmetry is unfolding just as AI's own architects concede they can no longer fully predict or control what they are building. A Global South with no visibility into that risk, no leverage over it, and no pooled mechanism to demand a role in shaping it, is left exposed twice over: first to dangers it did not define, and second to a governance order it had no hand in writing.

Feeding the Machine, Absent from the Rule book

In 2021, OpenAI contracted a San Francisco outsourcing firm called Sama to hire workers in Nairobi to read and label some of the most disturbing text on the internet such as, descriptions of child sexual abuse, bestiality, and graphic violence — so that ChatGPT's public release could later filter it out. Workers were reportedly paid between \$1.32 and \$2 an hour.⁵ That labor sits inside a product now used by hundreds of millions of people, including a fast-growing share of them in South Asia. The workers who did the labeling were never in the room when the resulting safety policies were written, and neither was their government.

This is the shape of the relationship, repeated across the industry in different forms. Content moderation for major platforms has for years relied on workers in Manila reviewing graphic and violent material for a fraction of what the equivalent work commands in the United States, often without adequate psychological support.⁶ Training data is harvested at scale from the online activity of users in Lagos, Jakarta, and Chittagong. None of that participation converts into a seat at the tables where the technology's actual rules are written — and those tables are easy to name, because there are so few of them and the guest list barely changes.

The UN Secretary-General's High-Level Advisory Body on AI released its final report, *Governing AI for Humanity*⁷, on 19 September 2024, calling for more inclusive global coordination — a call made necessary by the fact that the process producing the report was dominated by government and industry voices from the United States, China, and Europe. The UK's AI Safety Summit at Bletchley Park in November 2023 produced a declaration signed by 28 countries and the European

⁵Billy Perrigo, "OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic," *Time*, January 18, 2023.

⁶"Content Moderators at YouTube, Facebook and Twitter See the Worst of the Web — and Suffer Silently," *Washington Post*, July 25, 2019.

⁷United Nations, *Governing AI for Humanity: Final Report*, High-Level Advisory Body on Artificial Intelligence, September 19, 2024.

Union; Bangladesh, Pakistan, Sri Lanka, and Nepal were not among the signatories, and India stood in for the entire region.⁸

The Seoul Summit that followed in May 2024 narrowed the circle further still: only eleven governments signed the resulting declaration, nearly all of them Western economies plus Japan and the Republic of Korea, and South Asia had no signatory at all.⁹ The Paris AI Action Summit in February 2025 broke that pattern, but not for South Asia as a region — it broke it for India alone, which co-chaired the summit jointly with France and has since secured hosting rights for the 2026 AI Impact Summit in New Delhi.¹⁰ Three summits in, the pattern holds: when South Asia is in the room, it is India, by itself.

The G7's Hiroshima AI Process, launched in 2023, produced a voluntary code of conduct for what it called 'advanced AI systems' — drafted by seven of the world's wealthiest economies, none of them in the Global South, for a technology whose consequences will fall hardest outside that group.¹¹

A Failure to Organize

It would be convenient to end the argument there: a handful of rich, largely Western governments wrote the rules and left everyone else out. But the picture is not that simple, and pretending otherwise lets Global South governments off a hook they deserve to be on.

The European Union did not get a seat at the table by being invited. It built the table. The EU AI Act, which entered into force in August 2024, is the product of years of negotiation among 27 member states that disagree on a great deal, but chose to negotiate as a bloc precisely because none of them could individually have shaped global AI rules alone.¹² Singapore, a country of roughly

⁸UK Government, "The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023," GOV.UK. Signatories: Australia, Brazil, Canada, Chile, China, the EU, France, Germany, India, Indonesia, Ireland, Israel, Italy, Japan, Kenya, Saudi Arabia, the Netherlands, Nigeria, the Philippines, South Korea, Rwanda, Singapore, Spain, Switzerland, Türkiye, Ukraine, the UAE, the UK, and the US.

⁹"The Seoul Declaration by Countries Attending the AI Seoul Summit, 21–22 May 2024," signed by Australia, Canada, the EU, France, Germany, Italy, Japan, South Korea, Singapore, the UK, and the US.


¹⁰AI Action Summit 2025 (Paris, 10–11 February 2025), co-chaired by President Macron and Prime Minister Modi; the 2026 AI Impact Summit is scheduled for New Delhi, India, with the 2027 summit set for Geneva.

¹¹G7, "Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems," October 30, 2023.

¹²European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), published in the Official Journal July 12, 2024, entered into force August 1, 2024.

six million people, published a Model AI Governance Framework in January 2019 and later built an AI testing toolkit, AI Verify, launched in May 2022, that other governments now reference as a template.¹³ The United Arab Emirates appointed the world's first Minister of State for Artificial Intelligence in October 2017 and has since used sovereign investment to buy itself a stake inside the frontier AI industry, including a \$1.5 billion Microsoft investment in the Emirati firm G42 announced in April 2024.¹⁴ None of these states approaches South Asia's population, market size, or share of the internet's next billion users. What they share is a decision to act as a coordinated actor rather than a collection of individually weak ones.

A Decade of Individual and Bloc AI Governance Instruments



South Asia is not the only part of the Global South facing this choice, and it is not the only one that has chosen differently. The African Union adopted a Continental Artificial Intelligence Strategy in July 2024, giving 55 member states — most of them individually far poorer and more compute-constrained than India — a single document to point to in international negotiations and a mandate for a common position on data governance and AI investment.¹⁵ Whatever the strategy's limitations in implementation, its existence means African governments arrive at the next summit with something South Asian governments do not: a text that a regional body agreed to jointly, and

¹³Infocomm Media Development Authority of Singapore, "Model AI Governance Framework," first released January 2019, updated 2020; "AI Verify," launched as a minimum viable product May 2022.

¹⁴Omar Al Olama appointed UAE Minister of State for Artificial Intelligence, October 2017; Microsoft's \$1.5 billion investment in G42 announced April 2024.

¹⁵African Union, Continental Artificial Intelligence Strategy, adopted by the AU Executive Council, Accra, Ghana, July 18–19, 2024.

a claim that a slight to one member is a slight to the bloc. South Asia, with a combined population larger than Africa's and a considerably more mature IT services sector, has produced no equivalent.

By contrast, South Asia has no equivalent regional instrument. SAARC has been largely dormant as a forum for new policy coordination for close to a decade; BIMSTEC has prioritized connectivity and security cooperation, with technology discussions concentrated on cybersecurity rather than AI governance specifically. Individual states have moved on their own — India's NITI Aayog published a National Strategy for Artificial Intelligence in June 2018, and Bangladesh published its own National Strategy for Artificial Intelligence in March 2020, followed by a National AI Policy that remained in draft form through 2024 and into 2026 — but a national strategy is a domestic document, not a seat at an international summit, and none of it has yet translated into coordinated regional diplomacy at Bletchley, Seoul, or Paris.¹⁶

This matters because the excuse of powerlessness does not survive contact with India's own record. India is not a small state without leverage: it holds one of the world's largest pools of AI engineering talent, a domestic technology industry that staffs a meaningful share of Silicon Valley's own workforce, and has gone from a single signature at Bletchley, to co-chairing the summit in Paris, to hosting the one that follows it.¹⁷ If a South Asian state has already earned entry to the room, the absence of a coordinated regional position standing behind it is a missed opportunity, not an inevitability.

Part of the reason no South Asian government has pushed hard for that seat is that unilateral leverage is genuinely thin. A single country moving to regulate frontier AI on its own terms risks a response the industry has already shown it is willing to give: withdrawal, delayed access to new models, or compliance costs passed on to users rather than absorbed by the company. Italy's suspension of ChatGPT in March 2023 over data-protection concerns produced roughly a month of negotiation with OpenAI before service resumed on the regulator's terms, in late April — a negotiation Italy could hold because it was backed by the weight of the EU's data-protection regime


¹⁶NITI Aayog, National Strategy for Artificial Intelligence, Government of India, June 2018; Bangladesh, National Strategy for Artificial Intelligence, Information and Communication Technology Division, March 2020; National Artificial Intelligence Policy (draft), 2024, with a further draft under public consultation through February 2026.

¹⁷See notes 7 and 9 above.

behind it.¹⁸ A South Asian state attempting the same maneuver alone, without a comparable regional instrument to invoke, has far less certainty the outcome would end the same way. That is not a reason to stay silent. It is the strongest argument for building the collective instrument in the first place.

Whose Language Counts as Default

Even if the political absence described above were resolved tomorrow, a second and less visible exclusion would remain, built into the systems themselves.



Source: Akamai¹⁹

A widely cited 2020 study of linguistic diversity in natural language processing sorted the world's roughly 7,000 languages into six resource classes, based on how much labeled and unlabeled data each has available. The top class — what the study's authors call 'the Winners' — contains just seven languages, led by English, Spanish, German, Japanese, and French, together commanding a wildly disproportionate share of NLP research investment.²⁰ Bengali, spoken by more than 230 million people and the first language of both Bangladesh and a large part of West Bengal, sits two tiers below it, among languages the same authors call 'the Rising Stars' — a strong online presence

¹⁸Garante per la Protezione dei Dati Personali, emergency order against OpenAI, March 30, 2023; conditional lifting of the order, April 11, 2023, with core compliance measures due April 30 and service restored April 28, 2023.

¹⁹ <https://www.akamai.com/glossary/natural-language-processing>

²⁰Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury, "The State and Fate of Linguistic Diversity and Inclusion in the NLP World," Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (2020): 6282–6293.

undercut by a persistent shortage of labeled data. Nepali and Sinhala fare far worse still, closer to the bottom of the taxonomy than the top.

This is not a minor technical footnote. It means the default behavior of a chatbot — its cultural reference points, the political and religious sensitivities it is calibrated to avoid, the assumptions built into its content moderation — is set first for an English-speaking, largely Western user, and adjusted for a South Asian user only as an afterthought, where it is adjusted at all. A Bengali-language question on a locally sensitive religious or political matter is more likely to draw on secondhand, poorly represented patterns than one asked in English about an American context. The region does not train the models that shape it; it receives them, already formed.

The consequences of that gap are not abstract, and Bangladesh does not have to look far to see them. The United Nations' own fact-finding mission on Myanmar concluded in 2018 that Facebook had been 'a useful instrument' for those spreading hatred against the Rohingya, pointing in part to the platform's inability to moderate Burmese-language content at the scale and speed the crisis demanded.²¹ More than a million Rohingya refugees now live in camps around Cox's Bazar, a displacement the UN investigation linked in part to exactly the kind of failure this section describes: a platform built and safety-tuned around a small number of high-resource languages, deployed into a linguistic and political context its own systems could not adequately read. The lesson was not confined to Myanmar in 2017. It is a preview of what happens when a technology's safety infrastructure is calibrated for languages other than the ones spoken where the technology is actually deployed.

Political absence from AI governance means South Asia has no vote on the rules. Epistemic absence from the training process means it has no vote on the underlying assumptions — and assumptions, unlike rules, are not written down anywhere that can later be renegotiated at a summit. They are compressed into billions of parameters and shipped as a finished product.

²¹UN Human Rights Council, Report of the Independent International Fact-Finding Mission on Myanmar, A/HRC/39/64, 2018.

What a Seat Would Actually Require

Naming the problem is the easy part. Most commentary on AI and the Global South arrives, eventually, at a call for 'greater inclusion' or a 'stronger voice,' and then stops — because inclusion is not something the states currently benefiting from its absence hand out voluntarily. It has to be built, and there is a working precedent for how.

The Alliance of Small Island States, known as AOSIS, represents 39 countries with a combined population smaller than Bangladesh's alone and a negligible share of global GDP. Individually, none of its members could move a single line of an international climate agreement. Collectively, AOSIS is widely credited with pushing the 1.5-degree Celsius target into the Paris Agreement and, thirty years after first proposing the idea, with securing the creation of a dedicated Loss and Damage Fund at COP27 in November 2022 — an outcome wealthy states resisted for years, until a coordinated bloc made the cost of continued resistance higher than the cost of concession.²² AOSIS did not prevail by outweighing the states across the table from it. It prevailed by refusing to be negotiated with one small, individually dismissible government at a time.



Source: The Japan Times²³

²²UNFCCC, "COP27 Reaches Breakthrough Agreement on New 'Loss and Damage' Fund," November 20, 2022; Alliance of Small Island States, COP27 statements, November 2022 (39 member states).

²³ <https://www.japantimes.co.jp/business/2026/09/21/ai-markets-us-economy-risk/>

A Global South AI Governance Consortium built on the same logic would not need to begin as a treaty body or a new UN agency, either would take longer to build than the timeline allows. It would need three things that a coalition of BIMSTEC, ASEAN, and African Union member states could plausibly assemble within a single summit cycle: a shared technical secretariat capable of auditing frontier models independently for the risks South Asian and African societies actually face, rather than inheriting safety evaluations designed around American and European harms; a single negotiating position on data governance and compute access, so that Bangladesh is not left bargaining alone with a company valued in the hundreds of billions of dollars; and a standing seat, held collectively rather than by any one government, at whichever summit follows the one India is hosting in 2026. None of this requires the consortium's members to agree on much else. AOSIS's members do not agree on trade, fisheries policy, or their respective relationships with major powers; they organize around one shared exposure, rising seas, and nothing more. A Global South AI consortium would not need its members to agree on data-privacy philosophy or the proper size of government. They already share one exposure: dependence on systems they did not build, cannot audit, and had no hand in constraining.

The obvious objection is that frontier AI companies do not need Global South revenue the way oil-importing economies once needed OPEC's crude, and so a consortium modeled on collective bargaining has nothing real to withhold. That objection undersells the numbers. The users who will determine whether the next generation of AI products succeeds commercially are increasingly concentrated outside the wealthy economies that currently govern the technology — a Global South that, in South Asia alone, is home to a quarter of humanity. A bloc that can set common terms for data licensing, local deployment, and compliance with jointly agreed safety audits has leverage that individual governments negotiating one country at a time do not. The leverage is not in what the Global South can withhold today. It is in what frontier labs will need tomorrow, and in whether the terms for that access are set now, while there is still a negotiation to be had, or later, once dependence is already built into the product.

The window for building this leverage is not indefinite. Compute costs are falling, capability curves are steepening, and the gap between what frontier labs can build and what any outside body, inside or outside government, can verify is, by the admission of the people building it, widening rather than closing. A consortium assembled after the next generation of frontier systems is already

deployed will be negotiating over a fait accompli, the way climate diplomacy now negotiates over emissions already sitting in the atmosphere. The alternative to organizing now is not a neutral default; it is accepting extraction and exposure as a permanent condition, on the theory that someone else's safety summit will eventually get around to what happens in Dhaka.

Bangladesh will not decide whether the next frontier model is trained safely. It cannot, on its own, and no realistic amount of national policy work changes that fact. What it can help decide, alongside states that share its exposure and little else, is whether it accepts the terms it is given or spends the years before the next summit building the instrument that lets it negotiate them. Extraction without agency is not a law of nature. It is a status quo that persists for exactly as long as no one organizes against it.